The pdfTeX user manual

Hàn Thế Thành and others

pdfTeX 1.40.25

January 7, 2024

`https://pdftex.org`

Contents

Chapter 1

Introduction

The main purpose of the pdfTeX project is to create and maintain an extension of TeX that can produce PDF output directly from TeX source files and improve/enhance the result of TeX typesetting with the help of PDF output. When PDF output is not selected, pdfTeX produces standard DVI output. An important aspect of the project is to investigate alternative justification algorithms; notably, the margin kerning and font expansion algorithms following the HZ microtypography algorithm by Prof. Hermann Zapf.

pdfTeX is maintained by Hàn Thế Thành, the original author, and others. The pdfTeX home page is `https://www.pdftex.org`. Please send bug reports, suggestions, etc., to the mailing list (`https://lists.tug.org/pdftex`).

pdfTeX is based on the original TeX sources and Web2C, and has been successfully compiled on Unix, Windows and many other systems. It is actively maintained, with great care taken to keep new pdfTeX versions backward-compatible with earlier ones.

A conservative successor to TeX, named ε-TeX, was developed in the 1990s. Since pdfTeX version 1.40, the ε-TeX extensions are compiled as part of the pdfTeX engine and thus always available. For documentation on the ε-TeX extensions, see `https://ctan.org/pkg/etex`.

Furthermore, pdfTeX itself has acquired plenty of extensions over the years which are not related specifically to PDF output, generally new primitives for various features that are inconvenient or impossible to implement at the TeX level. Many of these extensions have been adopted across all engines (sometimes with different primitive names or variant functionality), and some are required by LaTeX. Therefore, in most distributions `etex` is a link to `pdftex`; the difference being whether DVI or PDF output is generated by default.

Other extensions are mlTeX and encTeX; these are also included in the pdfTeX implementation, although they are rarely used for new documents.

pdfTeX does not natively support UTF-8 input text, Unicode-encoded fonts, or anything else related to Unicode, as it was written long before Unicode became widespread. Making those changes to the engine now would necessarily create unacceptable incompatibilities, so there are no plans to do so. Thus, when using pdfTeX, LaTeX and other formats handle UTF-8 (and other) input at the TeX macro level, which works well enough in practice for most documents. It is also possible to use TrueType and OpenType fonts with pdfTeX, if you choose an 8-bit subset to be encoded.

If you need a TeX engine with native support for Unicode input, TrueType fonts, OpenType fonts, please look into LuaTeX (`https://ctan.org/pkg/luatex`) or XeTeX (`https://tug.org/xetex`).

1.1 About this manual

This manual revision (909) covers pdfTeX development up to version 1.40.25. The primary repository for both the manual and the pdfTeX sources is `svn://tug.org/pdftex/branches/stable`. The typeset manual in PDF format can be found on CTAN in `https://ctan.org/pkg/pdftex`.

Thanks to the many people who have contributed to the manual. Improvements are always possible, and bugs not unlikely. Please send questions or suggestions via email at `https://lists.tug.org/pdftex`.

1.2 Legal notice

Copyright © 1996–2024 Hàn Thế Thành. Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled "GNU Free Documentation License".

1.3 About PDF

Let's start with a very brief introduction to PDF format. For our example, the bit of TeX source below (figure 1.1) generates the nearly-minimal PDF file shown on the next page (figure 1.2). Since compression is disabled, such a PDF file is rather verbose and readable. The first line (`%PDF-1.4`) specifies the PDF version used. PDF viewers are supposed to silently skip over all elements they cannot handle.

A PDF file consists of objects. These objects can be recognized by their number and keywords. For example, over in the second column, we can see (modulo editorial line breaks):

```
9 0 obj << /Type /Catalog /Pages 5 0 R >> endobj
```

Here `9 0 obj ... endobj` is the object capsule. The first number is the object number (no. 9). The sequence `5 0 R` is an object reference, that is, a pointer to another object (no. 5). The second number (here a zero) is currently not used in pdfTeX; it is the version number of the object. It is, for instance, used by PDF editors, when they replace objects by new ones.

When a viewer opens a PDF file, it goes to the end of the file, looking for the keyword `startxref`. The number after `startxref` gives the absolute position (byte offset from the file start) of the so-called "object cross-reference table" that begins with the keyword `xref`. This table in turn tells the byte offsets of all objects that make up the PDF file, providing fast random access to the individual

---

```
\pdfoutput=1
\pdfcompresslevel=0
\pdfobjcompresslevel=0
\pdfmapline{ptmr8r Times-Roman 2 <8r.enc}
\font\tenrm=ptmr8r
\tenrm
Welcome to pdf\TeX!
\end
```

Figure 1.1: This plain TeX source generates PDF output shown in figure 1.2.

```
%PDF-1.4                                           4 0 obj
%                                                  <<
3 0 obj                                            /Type /Font
<<                                                 /Subtype /Type1
/Length 119                                        /BaseFont /Times-Roman
>>                                                 /FontDescriptor 8 0 R
stream                                             /FirstChar 33
BT                                                 /LastChar 116
/F1 9.9626 Tf 72 713.245 Td                        /Widths 7 0 R
[(W)80(elcome)-250(to)-250(pdfT)]TJ 67.818         /Encoding 6 0 R
-2.241 Td [(E)]TJ 4.842 2.241 Td [(X!)]TJ          >>
ET                                                 endobj
endstream                                          5 0 obj
endobj                                             <<
2 0 obj                                            /Type /Pages
<<                                                 /Count 1
/Type /Page                                        /Kids [2 0 R]
/Contents 3 0 R                                    >>
/Resources 1 0 R                                   endobj
/MediaBox [0 0 612 792]                            9 0 obj
/Parent 5 0 R                                      <<
>>                                                 /Type /Catalog
endobj                                             /Pages 5 0 R
1 0 obj                                            >>
<<                                                 endobj
/Font << /F1 4 0 R >>                              10 0 obj
/ProcSet [ /PDF /Text ]                            <<
>>                                                 /Producer (pdfTeX-1.40.25)
endobj                                             /Creator (TeX)
7 0 obj                                            /CreationDate (D:20240106180300-08'00')
[333 408 500 500 833 778 333 333 333 500           /ModDate (D:20240106180300-08'00')
564                                                /Trapped /False
250 333 250 278 500 500 500 500 500 500 500        /PTEX.Fullbanner (This is pdfTeX, Version
500 500 500 278 278 564 564 564 444 921 722        3.141592653-2.6-1.40.25 (TeX Live 2023)
667 667 722 611 556 722 722 333 389 722 611        kpathsea version 6.3.5)
889 722 722 556 722 667 556 611 722 722 944        >>
722 722 611 333 278 333 469 500 333 444 500        endobj
444 500 444 333 500 500 278 278 500 278 778        xref
500 500 500 500 333 389 278]                       0 11
endobj                                             0000000000 65535 f
8 0 obj                                            0000000296 00000 n
<<                                                 0000000192 00000 n
/Type /FontDescriptor                              0000000015 00000 n
/FontName /Times-Roman                             0000001010 00000 n
/Flags 2                                           0000001162 00000 n
/FontBBox [0 -216 1000 678]                        0000000898 00000 n
/Ascent 678                                        0000000363 00000 n
/CapHeight 651                                     0000000716 00000 n
/Descent -216                                      0000001219 00000 n
/ItalicAngle 0                                     0000001268 00000 n
/StemV 83                                          trailer
/XHeight 450                                       << /Size 11
>>                                                 /Root 9 0 R
endobj                                             /Info 10 0 R
6 0 obj                                            /ID [<BFF45DC9727C079A3E1AD02B993FCE84>
<<                                                 <BFF45DC9727C079A3E1AD02B993FCE84>] >>
/Type /Encoding                                    startxref
/Differences [33/exclam 69/E 84/T 87/W/X           1529
99/c/d/e/f 108/l/m 111/o/p 116/t]                  %%EOF
>>
endobj
```

Figure 1.2: The PDF output for the TEX source in figure 1.1.

objects (here the `xref` table shows 11 objects, numbered from 0 to 10; object no. 0 is always unused). The actual starting point of the file's object structure is defined after the `trailer`: the `/Root` entry points to the `/Catalog` object (no. 9). In this object the viewer can find the pointer `/Pages` to the page list object (no. 5). In our example we have only one page. The trailer also usually holds an `/Info` entry, which points to an object (no. 10) with a bit more about the document. We can follow the thread:

`/Root` ⟶ object 9 ⟶ `/Pages` ⟶ object 5 ⟶ `/Kids` ⟶ object 2 ⟶ `/Contents` ⟶ object 3

As soon as we add annotations, a fancy word for hyperlinks and the like, some more entries will be present in the catalog. We invite users to take a look at the PDF code of this file to get an impression of that.

The page content is a stream of drawing operations. Such a stream can be compressed, where the level of compression can be set with `\pdfcompresslevel` (compression is switched off for the title page). Let's take a closer look at this stream in object 3. Often (but not in our example) there is a transformation matrix, six numbers followed by `cm`. As in PostScript, the operator comes after the operands. Between `BT` and `ET` comes the text. A font is selected by a `Tf` operator, which is given a font resource name `/F..` and the font size. The actual text goes into parentheses `()` so that it creates a PostScript string. The numbers in bracket pairs provide horizontal movements like spaces and fine glyph positioning (kerning). When one analyzes a file produced by a less sophisticated typesetting engine, whole sequences of words can be recognized. In PDF files generated by pdfTeX however, many words come out rather fragmented, mainly because a lot of kerning takes place; in our example the `80` moves the text `(elcome)` left towards the letter `(W)` by $80/1000$ of the font size. PDF viewers in search mode simply ignore the kerning information in these text streams. When a document is searched, the search engine reconstructs the text from these (string) snippets.

Every `/Page` object points also to a `/Resources` object (no. 1) that gives all ingredients needed to assemble the page. In our example only a `/Font` object (no. 4) is referenced, which in turn tells that the text is typeset in `/Font /Times-Roman`. The `/Font` object points also to a `/Widths` array (object no. 7) that tells for each character by how much the viewer must move forward horizontally after typesetting a glyph.

More details about the font can be found in the `/FontDescriptor` object (no. 8); if a font file is embedded, this object points to the font program stream. But as the Times-Roman font used for our example is one of the 14 so-called standard fonts that should always be present in any PDF viewer and therefore need not be embedded in the PDF file, it is left out here for brevity. However, when we use for instance a Computer Modern Roman font, we have to make sure that this font is later available to the PDF viewer, and the best way to do this is to embed the font. It's highly recommended nowadays to embed even the standard fonts; you can't know how it looks exactly at the viewer side unless you embed every font.

In this simple file we don't specify in what way the file should be opened, for instance full screen or clipped. A closer look at the page object no. 2 (`/Type/Page`) shows that a mediabox (`/MediaBox`) is part of the page description. A mediabox acts like the (high-resolution) bounding box in a PostScript file. pdfTeX users can add dictionary entries to page objects with the `\pdfpageattr` primitive.

Although in most cases macro packages will shield users from these internals, pdfTeX provides access to many of the entries described here, either automatically by translating the TeX data structures into PDF ones, or manually by pushing entries to the catalog, page, info or self-created objects. One can for instance create an object by using `\pdfobj`, after which `\pdflastobj` returns its number. So

```
\pdfobj { << /Type/ExtGState /LW 2 >> }
```

inserts an object into the PDF file (it creates a "graphics state" object setting the line width to 2 units), and `\pdflastobj` now returns the number pdfTEX assigned to this object. Unless objects are referenced by others, they will just end up as isolated entities, not doing any real harm but bloating the PDF file.

In general this rather direct way of pushing objects in the PDF files by primitives like `\pdfobj` is not very useful, and only makes sense when implementing, say, fill-in field support or annotation content reuse. We will come to that later.

Of course, this is just the barest introduction to PDF format. For those who want to learn more about technical PDF details, the best bet is to read the PDF reference manual (`https://pdfa.org/resource/pdf-specification-index/`).

We turn now to the specifics of pdfTEX.

Chapter 2

Invoking pdfTeX

pdfTeX has many command line options and can use many environment variables and configuration file settings. Except for the simple and rarely-used `-draftmode` and `-output-format` options, they are all inherited from the common framework for TeX engines as implemented in Web2C and Kpathsea. Their documentation is available at `https://tug.org/web2c` and `https://tug.org/kpathsea`.

Two additional environment variables need more description here: first, `SOURCE_DATE_EPOCH` (introduced in version 1.40.17, 2016). If this is set, it must be a positive integer (with one trivial exception: if it is set but empty, that is equivalent to 0). Non-integer values cause a fatal error. The value is used as the current time in seconds since the usual Unix "epoch": the beginning of 1970-01-01, UTC. Thus, a value of `32` would result in a `/CreationDate` and `/ModDate` values of `19700101000032Z`. This is useful for reproducible builds of documents. (See also `\pdfinfoomitdate`, `\pdfsuppressptexinfo`, et al.)

The second, related, environment variable is `FORCE_SOURCE_DATE`. If this is set to 1, TeX's time-related primitives are also initialized from the value of `SOURCE_DATE_EPOCH`. These primitives are `\year`, `\month`, `\day`, and `\time`. If `SOURCE_DATE_EPOCH` is not set, setting `FORCE_SOURCE_DATE` has no effect. If `FORCE_SOURCE_DATE` is unset, set to the empty string, or set to `0`, the primitives reflect the current time as usual. Any other value elicits a warning, and the current time is used. This is useful if one wants to make reproducible PDFs for a set of documents without changing them in any way, e.g., an operating system distribution with manuals that use `\today`. Except in such unusual circumstances, it is better not to set this, and let the TeX primitives retain the meaning they have always had.

In addition, if both `SOURCE_DATE_EPOCH` and `FORCE_SOURCE_DATE` are set, `\pdffilemoddate` returns a value in UTC, ending in Z. (The values of the environment variables are irrelevant in this case.)

Finally, just to have the list of options and basic invocation at hand, here is a verbatim listing of the `--help` and `--version` output. All options can be specified with one or two dashes and unambiguously abbreviated.

```
Usage: pdftex [OPTION]... [TEXNAME[.tex]] [COMMANDS]
   or: pdftex [OPTION]... \FIRST-LINE
   or: pdftex [OPTION]... &FMT ARGS
  Run pdfTeX on TEXNAME, usually creating TEXNAME.pdf.
  Any remaining COMMANDS are processed as pdfTeX input, after TEXNAME is read.
  If the first line of TEXNAME is %&FMT, and FMT is an existing .fmt file,
  use it.  Else use 'NAME.fmt', where NAME is the program invocation name,
  most commonly 'pdftex'.
```

Alternatively, if the first non-option argument begins with a backslash,
interpret all non-option arguments as a line of pdfTeX input.

Alternatively, if the first non-option argument begins with a &, the
next word is taken as the FMT to read, overriding all else. Any
remaining arguments are processed as above.

If no arguments or options are specified, prompt for input.

```
-cnf-line=STRING        parse STRING as a configuration file line
-draftmode              switch on draft mode (generates no output PDF)
-enc                    enable encTeX extensions such as \mubyte
-etex                   enable e-TeX extensions
[-no]-file-line-error   disable/enable file:line:error style messages
-fmt=FMTNAME            use FMTNAME instead of program name or a %& line
-halt-on-error          stop processing at the first error
-ini                    be pdfinitex, for dumping formats; this is implicitly
                          true if the program name is 'pdfinitex'
-interaction=STRING     set interaction mode (STRING=batchmode/nonstopmode/
                          scrollmode/errorstopmode)
-ipc                    send DVI output to a socket as well as the usual
                          output file
-ipc-start              as -ipc, and also start the server at the other end
-jobname=STRING         set the job name to STRING
-kpathsea-debug=NUMBER  set path searching debugging flags according to
                          the bits of NUMBER
[-no]-mktex=FMT         disable/enable mktexFMT generation (FMT=tex/tfm/pk)
-mltex                  enable MLTeX extensions such as \charsubdef
-output-comment=STRING  use STRING for DVI file comment instead of date
                          (no effect for PDF)
-output-directory=DIR   use existing DIR as the directory to write files in
-output-format=FORMAT   use FORMAT for job output; FORMAT is 'dvi' or 'pdf'
[-no]-parse-first-line  disable/enable parsing of first line of input file
-progname=STRING        set program (and fmt) name to STRING
-recorder               enable filename recorder
[-no]-shell-escape      disable/enable \write18{SHELL COMMAND}
-shell-restricted       enable restricted \write18
-src-specials           insert source specials into the DVI file
-src-specials=WHERE     insert source specials in certain places of
                          the DVI file. WHERE is a comma-separated value
                          list: cr display hbox math par parend vbox
-synctex=NUMBER         generate SyncTeX data for previewers according to
                          bits of NUMBER ('man synctex' for details)
-translate-file=TCXNAME use the TCX file TCXNAME
-8bit                   make all characters printable by default
-help                   display this help and exit
-version                output version information and exit
```

pdfTeX home page: <http://pdftex.org>

Email bug reports to pdftex@tug.org (https://lists.tug.org/pdftex).

pdfTeX 3.141592653-2.6-1.40.25 (TeX Live 2023)

```
kpathsea version 6.3.5
Copyright 2023 Han The Thanh (pdfTeX) et al.
There is NO warranty.  Redistribution of this software is
covered by the terms of both the pdfTeX copyright and
the Lesser GNU General Public License.
For more information about these matters, see the file
named COPYING and the pdfTeX source.
Primary author of pdfTeX: Han The Thanh (pdfTeX) et al.
Compiled with libpng 1.6.39; using libpng 1.6.39
Compiled with zlib 1.2.13; using zlib 1.2.13
Compiled with xpdf version 4.04
```

2.1   Macro packages supporting pdfTEX

Currently all mainstream macro packages offer pdfTEX support, with automatic detection of pdfTEX as engine. So there is normally no need to turn on pdfTEX support explicitly.

- For LATEX users, the `hyperref` package (originally written by Sebastian Rahtz and Heiko Oberdiek; now maintained by the LATEX team, has substantial support for pdfTEX and provides access to most of its features. In the simplest and most common case, the user merely needs to load `hyperref`, and all cross-references will be converted to PDF hypertext links. PDF output is automatically selected, compression is turned on, and the PDF page size is set up correctly. Bookmarks are created to match the table of contents.

- The standard LATEX packages `graphics`, `graphicx`, and `color` also have built-in pdfTEX support, which allows use of color, text rotation, and graphics inclusion commands.

- The ConTEXt MkII system by Hans Hagen has full support for pdfTEX in its generalized hypertext features. The latest ConTEXt supports advanced PDF output, but uses a engine (LMTX).

- PDF from GNU Texinfo documents can be created by running pdfTEX on the Texinfo file, instead of TEX. Alternatively, run the shell command `texi2pdf` instead of `texi2dvi`.

- For plain TEX users, the `miniltx.tex` file from the `graphics-pln` package allows loading `graphics`, `graphicx`, and `color`. Eplain provides additional support for hyperlinks.

- A modification of `webmac.tex`, named `pdfwebmac.tex`, allows production of hyperlinked PDF versions of the literate source code written in WEB, such as pdfTEX.

For TEX developers: as pdfTEX generates the final PDF output without help of a postprocessor, macro packages that take care of these PDF features have to be set up properly. Tasks include handling color, graphics, hyperlink support, threading, fonts, page imposition and manipulation. All of these PDF-specific tasks can be controlled by pdfTEX's own primitives (a few also by a pdfTEX-specific `\special{pdf: ...}`). Any other `\special` commands, like the ones defined for various DVI postprocessors, are simply ignored by pdfTEX when in PDF output mode; a warning is given for non-empty `\special` commands.

When a macro package written for classical TEX with DVI output is to be modified for use with pdfTEX, it is helpful to get some insight to what extent pdfTEX-specific support is needed. This info can be gathered, for instance, by outputting the various `\special` commands via `\message`. as in:

```
\pdfoutput=1 \let\special\message
```

or, if this leads to confusion,

```
\pdfoutput=1 \def\special#1{\write16{special: #1}}
```

and see what happens. As soon as one `special:` message turns up, one knows for sure that some kind of pdfTEX-specific support is needed, and often the message itself gives a indication of what is needed.

Chapter 3

Setting up fonts

pdfTeX can work with Type 1 and TrueType fonts (and to a small extent also with OpenType fonts). Font files should be available and embedded for all fonts used in the generated PDF. It is possible to use METAFONT-generated fonts in pdfTeX—but it is strongly recommended not to use these fonts if an equivalent is available in Type 1 or TrueType format, if only because bitmap Type 3 fonts render poorly under enlargement.

3.1 Map files

Font map files provide the connection between TeX TFM font files and outline font file names. They contain also information about re-encoding arrays, partial font embedding ("subsetting"), and character transformation parameters (like SlantFont and ExtendFont). Those map files were first created for DVI postprocessors. But, as pdfTeX in PDF output mode includes all PDF processing steps, it also needs to know about font mapping, and therefore reads in one or more map files. Map files are not read in when pdfTeX is in DVI mode. Bitmap fonts can (and normally should) be used without being listed in the map file.

By default, pdfTeX reads the map file `pdftex.map`. In Web2C, map files are searched for using the `TEXFONTMAPS` config file value and environment variable. By default, the current directory and various system directories are searched.

Within the map file, each font is listed on a single line. The syntax of each line is upward-compatible with `dvips` map files and can contain the following fields (some are optional; explanations follow):

> *tfmname psname fontflags special encodingfile fontfile*

Here are two (real-world) examples. All the values are explained in detail in the following.

```
cmr10 CMR10 <cmr10.pfb
ptmr8y NimbusRomNo9L-Regu " TeXnANSIEncoding ReEncodeFont " <texnansi.enc <utmr8a.pfb
```

It is mandatory that *tfmname* is the first field. If a *psname* is given, it must be the second field. Similarly if *fontflags* is given it must be the third field (if *psname* is present) or the second field (if *psname* is left out). The positions of *special*, *encodingfile*, and *fontfile* can be mixed.

3.1.1 Map lines: *tfmname*

The *tfmname* field specifies the name of the TFM file for a font—the file name given in a TeX `\font` command. This name must always be given, with no extension. Examples: `cmr10`, `ptmr8y`.

3.1.2 Map lines: *psname*

The *psname* field specifies the PostScript (or other outline) font name, as defined within the outline font file. Examples: `CMR10`, `NimbusRomNo9L-Regu`. It is highly recommended to use the *psname* field, but strictly speaking it is optional. It has two main uses.

First, when a PDF file is embedded by `\pdfximage`, the `/BaseFont` names in the font dictionaries of Type 1 and Type 1C (CFF) fonts from the embedded PDF file are checked against this *psname* field. If names match, the glyphs of that font will not be copied from the embedded PDF file, but instead a local font is opened, and all needed glyphs will be taken from the Type 1 font file that is mentioned in the map line (see *fontfile* below). By this collecting mechanism Type 1 glyphs can be shared between several embedded PDF files and with text that is typeset by pdfTEX, which helps keep the resulting PDF file size smaller, if many files with similar Type 1(C) fonts are embedded. Replacing Type 1 fonts from embedded PDF files requires that also a Type 1 font file name is in the *fontfile* field (see below).

Second, if a font file is not to be embedded into the PDF output (*fontfile* field missing), then the *psname* field will be copied to the `/BaseFont` and `/FontName` dictionary entries in the PDF file, so that the PostScript font name will be known to viewers and other PDF-reading applications.

3.1.3 Map lines: *fontflags*

The *fontflags* field optionally specifies various characteristics of the font. The following description of these flags is taken, with slight modification, from the PDF reference manual (the section on font descriptor flags). Viewers may adapt their rendering to these flags, especially when they substitute for a non-embedded font.

The value of the flags key in a font descriptor is a 32-bit integer that contains a collection of boolean attributes. These attributes are true if the corresponding bit is set to 1. The following table specifies the meanings of the bits, with bit 1 being the least significant. Reserved bits must be set to zero.

bit position	semantics
1	Fixed-width font
2	Serif font
3	Symbolic font
4	Script font
5	Reserved
6	Uses the Adobe Standard Roman character set
7	Italic
8–16	Reserved
17	All-capitals font
18	Small-capitals font
19	Force bold at small text sizes
20–32	Reserved

The first several bits specify the general type of font. All characters in a *fixed-width* font (a.k.a. monospaced, typewriter) have the same width, while characters in a proportional font have different widths. Characters in a *serif font* have short strokes drawn at an angle on the top and bottom of character stems, while sans serif fonts do not have such strokes. A *symbolic font* contains symbols rather than letters and numbers. Characters in a *script font* resemble cursive handwriting. An *all-capitals* font, which is typically used for display purposes such as titles or headlines, contains no lowercase letters. It differs from a *small-capitals* font in that characters in the latter, while also capital letters, have been sized and their proportions adjusted so that they have the same size and stroke weight as lowercase characters in the same typeface family.

Bit 6 in the flags field indicates that the font's character set conforms to the Adobe Standard Roman Character Set, or a subset of that, and that it uses the standard names for those characters.

Finally, bit 19 is used to determine whether or not bold characters are drawn with extra pixels even at very small text sizes. Typically, when characters are drawn at small sizes on very low resolution devices such as display screens, features of bold characters may appear only one pixel wide. Because this is the minimum feature width on a pixel-based device, ordinary non-bold characters also appear with one-pixel wide features, and thus cannot be distinguished from bold characters. If bit 19 is set, features of bold characters may be thickened at small text sizes.

If the *fontflags* field is not given, and the font is embedded, pdfTeX treats it as the value 4 (decimal, that is, bit position 3 is set), a symbolic font. For non-embedded fonts, the default value is `0x22`, a non-symbolic serif font. If you do not know the correct value, it is best not to specify it at all, as specifying a bad value of font flags may cause trouble in viewers. On the other hand this option is not absolutely useless because it provides backward compatibility with older map files (see the *fontfile* description below).

3.1.4 Map lines: *special*

The *special* field specifies font manipulations in the same way as `dvips`. Currently only the keywords `SlantFont` and `ExtendFont` are interpreted; other instructions (notably `ReEncodeFont` and its parameters, see *encoding* below) are ignored. The permitted `SlantFont` range is $-1..1$; for `ExtendFont` it's $-2..2$. The text of the *special* field must be enclosed by ASCII double quote characters: `"`.

3.1.5 Map lines: *encodingfile*

The *encodingfile* field specifies the name of the file containing the external encoding vector to be used for the font. The encoding file name must have the extension `.enc`, and the file name including extension must be given with either a preceding `<` character or a preceding `<[`. The format of the encoding vector is identical to that used by `dvips`. If no encoding is specified, the font's built-in default encoding is used. The *encodingfile* field may be omitted if you are sure that the font resource has the correct built-in encoding. In general this option is highly recommended, and it is required when subsetting a TrueType font.

Starting with pdfTeX version 1.40.19, an encoding file can also be specified for bitmap PK fonts. In this case, it assigns the glyph names from the given encoding vector, which can be used with the `\pdfglyphtounicode` primitive (q.v.). For example:

```
\pdfglyphtounicode{ffi}{0066 0066 0069} % normally: \input glyphtounicode
\pdfgentounicode=1
\pdfmapline{cmb10 <7t.enc}
\font\cmb=cmb10 \cmb ffi
```

The result is a PDF file with a correctly-labeled `/ffi` character instead of the numeric character position in `cmb10.tfm` (decimal 14).

3.1.6 Map lines: *fontfile*

The *fontfile* field sets the name of the font file to be embedded into the PDF output for a given TeX font; the *tfmname* ↔ *fontfile* mapping is the most prominent use of the `pdftex.map` file.

The font file name must refer to a Type 1 or TrueType font file. If the *fontfile* field is missing, no font embedding can take place; a warning will be given, unless the *psname* field contain one of the 14 standard font names. Not embedding all fonts in a PDF file is troublesome, as this forces the PDF viewer to use or synthesize a replacement, typically with awful results.

The font file name should be preceded by one or two special characters, specifying how to handle the font file:

- If the font file name is preceded by a `<` character (as in `<cmr10.pfb`), the font file will be only partially embedded in the PDF output ("subsetted"), meaning that only used glyphs are written to the PDF file. This is the most common use and is *strongly recommended* for any font, as it ensures the portability and reduces the size of the PDF output. Subsetted fonts are included in such a way that name and cache clashes are minimized.

- If the font file name is preceded by a double `<<`, the font file will be included entirely—all glyphs of the font are embedded, including even those not used in the document. Apart from increasing the PDF output size, this option may cause troubles with TrueType fonts, so it is normally not recommended for Type 1 or TrueType fonts. But this is currently the only mode that allows the use of OpenType fonts. This mode might also be useful in case the font is atypical and cannot be subsetted well by pdfTeX. *Beware: proprietary font vendors typically forbid full font inclusion.*

- As of pdfTeX version 1.40.0, if no special character precedes the font file name, it is ignored, with a warning. You achieve exactly the same PDF result if you just remove the font file name from the map entry. Then the glyph widths that go into the PDF file are extracted from the TFM file, and a font descriptor object is created that contains approximations of the font metrics for the selected font.

- Specifying the *psname* and no font file name is only useful as a last-ditch fallback when you do not want to embed the font (e.g., due to font license restrictions), but wish to use the font metrics and let the PDF viewer generate instances that look close to the used font, in case the font resource is not installed on the system where the PDF output will be viewed or printed. To use this feature, the font flags *must* be specified, and it must have the bit 6 set on, which means that only fonts with the Adobe Standard Roman character set can be simulated. The only exception is the case of a symbolic font. In these days of Unicode, these font approximations are not likely to be useful.

If you encounter problematic lookups, for instance if pdfTeX tries to open a `.pfa` file instead of a `.pfb`, you can add the suffix to the filename. In this respect, pdfTeX completely relies on the Kpathsea library.

For Type 1 and TrueType fonts, the font file will be included only once in the PDF output, regardless of how many TeX `\font` instances are used in the document. For instance, given

```
\font\a = cmr12
\font\b = cmr12 at 11pt
```

the outline file `cmr12.pfb` will only be included once in the PDF, and merely scaled down to create the instance for `\b`.

If a used font is not present in the map files, pdfTeX will try to use PK fonts as most DVI drivers do, creating PK fonts on-the-fly if needed. This is the normal, and recommended, way to use bitmap fonts.

3.1.7   Map lines: summary

To summarize this rather complex story, let's look at some more example map lines. The most common way is to embed only a subset of glyphs from a font for a particular desired encoding, here `8r`:

```
ptmri8r Times-Italic <8r.enc <ptmri8a.pfb
```

Without re-encoding it looks like this:

```
cmr10 CMR10 <cmr10.pfb
```

`SlantFont` and `ExtendFont` fields are specified as with `dvips`:

```
psyro   StandardSymL ".167 SlantFont"          <usyr.pfb
pcrr8rn Courier      ".85 ExtendFont" <8r.enc <pcrr8a.pfb
```

Entirely embed a font into the PDF file without and with re-encoding (not typically useful):

```
fmvr8x MarVoSym          <<marvosym.pfb
pgsr8r GillSans <8r.enc <<pgsr8a.pfb
```

A TrueType font can be used in the same way as a Type 1 font:

```
verdana8r Verdana <8r.enc <verdana.ttf
```

Finally, a few cases with non-embedded fonts. If the font file is missing, the viewer application will have to use its own approximation of the missing font (with and without re-encoding):

```
ptmr8r Times-Roman <8r.enc
psyr Symbol
```

In the final example the numerical font flags (bit position 6) specify using the Adobe Standard Roman character set, so the viewer might try an approximation:

```
pgsr8r GillSans 32
```

Not embedding fonts is rather risky and should generally be avoided. The recommendation these days is to embed all fonts, even the 14 standard ones.

3.2   Helper tools for TrueType fonts: `ttf2afm`

As mentioned above, pdfTEX can work with TrueType fonts. Defining TrueType fonts is similar to Type 1. The only extra thing to do with TrueType is to create a TFM file. There is a program called `ttf2afm` in the pdfTEX distribution which can be used to extract AFM from TrueType fonts (another conversion program is `ttf2pt1`). Basic usage of `ttf2afm`:

```
ttf2afm -e encfile.enc -o output.afm input.ttf
```

A TrueType file can be recognized by its suffix `ttf`. If no `-o` option is given, `ttf2afm` writes the output AFM to standard output.

The optional *encfile* specifies the encoding, which is the same as the encoding vector used in map files for pdfTEX and `dvips`. That is, it must be an 8-bit encoding, not Unicode. If the encoding is not given, all the glyphs of the AFM output will be mapped to `/.notdef`. If we need to know which glyphs are available in the font, we can run `ttf2afm` without any `-e` to get all glyph names. The resulting AFM file can be used to generate a TFM by applying the `afm2tfm` utility.

To use a new TrueType font the minimal steps may look like below, supposing that a map file `test.map` is used.

```
ttf2afm -e 8r.enc -o times.afm times.ttf
afm2tfm times.afm -T 8r.enc
echo "times TimesNewRomanPSMT <8r.enc <times.ttf" >>test.map
```

There are a limitations with TrueType fonts in comparison with Type 1 fonts:

- The special effects SlantFont/ExtendFont did not work before version 1.40.0.

- To subset a TrueType font, the font must be specified as re-encoded, therefore an encoding vector must be given.

- TrueType fonts used in embedded PDF files are kept untouched; they are not replaced or merged with the same font used in the document, as happens with Type 1.

For much more about pdfTEX and TrueType fonts, including many details on handling glyph names, see "A closer look at TrueType fonts and pdfTEX", *TUGboat* 30:1 (2009), pp. 32-34, `https://tug.org/TUGboat/tb30-1/tb94thanh.pdf`.

Chapter 4

pdfTeX primitives

Here follows a description of the primitives added by pdfTeX to the original TeX engine (other extensions by ε-TeX, mlTeX and encTeX are not described). Many of these primitives are described further in the `samplepdf.tex` file in the pdfTeX distribution (q.v.).

If the output is DVI, then the pdfTeX-specific dimension parameters are not used at all. However, some pdfTeX integer parameters can affect DVI as well as PDF output (specifically, `\pdfoutput` and `\pdfadjustspacing`).

A warning to macro writers: if you define macros whose names start with `\pdf`, you risk name clashes with new primitives that may be introduced in future versions of pdfTeX.

General warning: many of these new primitives, for example `\pdfdest` and `\pdfoutline`, write their arguments directly to the PDF output file (when producing PDF), as PDF string constants. This means that *you* (or, more likely, the macros you write) must escape characters as necessary (namely `\`, `(`, and `)`. Otherwise, an invalid PDF file may result. The `hyperref` and Texinfo packages have code which may serve as a starting point for implementing this, although it will certainly need to be adapted to any particular situation.

4.1   Document setup

4.1.1   \pdfoutput

`\pdfoutput`   (integer)

This parameter specifies whether the output format should be DVI or PDF. A positive value means PDF output, otherwise (default 0) one gets DVI output. This primitive is the only one that must be set to produce PDF output (the command-line option `-output-format=pdf` may alternatively be used); all other primitives are optional. This parameter cannot be specified *after* shipping out the first page. In other words, to get PDF output, we have to set `\pdfoutput` before pdfTeX the first page.

A simple way of making macros aware of pdfTeX in both PDF or DVI mode is:

```
\ifx\pdfoutput\undefined \csname newcount\endcsname\pdfoutput \fi
\ifcase\pdfoutput DVI CODE \else PDF CODE \fi
```

Using the `ifpdf.sty` file, which works with both LaTeX and plain TeX, is a cleaner way of doing this. Originally, the simple test `\ifx\pdfoutput\undefined` sufficed; but for many years, the pdfTeX engine is used in distributions also for non-PDF formats (e.g., LaTeX), so `\pdfoutput` may be defined even when the output format is DVI.

4.1.2 \pdfmajorversion, \pdfminorversion

\pdfmajorversion   (integer)
\pdfminorversion   (integer)

Together, these two primitives specify the PDF version for generated PDF output. The defaults compiled into the pdfTeX program are \pdfmajorversion=1 and \pdfminorversion=4, thus PDF 1.4 is generated by default.

However, distributions typically alter the engine's compiled default minor version of 4 when building formats. For example, as of 2010, TeX Live sets \pdfminorversion=5 when formats are built. This is so object compression can be enabled (see \pdfobjcompress below).

This value also defines the highest PDF version for included PDFs to be allowed without error, by default (see \pdfinclusionerrorlevel).

The values for both must be ≥ 1 but are not checked further. Furthermore, they are set independently; setting only \pdfmajorversion=2 would result in PDF 2.4 output; it's necessary to additionally set \pdfminorversion.

If specified, these primitives must appear before any data is written to the generated PDF file. The \pdfmajorversion primitive was introduced in pdfTeX 1.40.21. \pdfminorversion was originally a synonym of the \pdfoptionpdfminorversion command, which is now obsolete. The primitive was introduced in pdfTeX 1.30.0.

4.1.3 \pdfcompresslevel

\pdfcompresslevel   (integer)

This integer parameter specifies the level of stream compression. Zero means no compression, 1 means fastest, 9 means best, 2..8 means something in between. A value outside this range will be adjusted to the nearest meaningful value. This parameter is read each time pdfTeX starts a stream.

This compression applies to text, inline graphics, and embedded PNG images (but only if they are un- and re-compressed during the embedding process). It is implemented by the zlib library).

Setting \pdfcompresslevel=0 is useful for PDF stream debugging.

4.1.4 \pdfobjcompresslevel

\pdfobjcompresslevel   (integer)

This integer parameter controls the compression of non-stream objects. If specified, the parameter must appear before any data is written to the PDF output. The primitive was introduced in pdfTeX 1.40.0.

In the PDF-1.4 specification, non-stream objects had to be written in the PDF file as clear text, uncompressed. The PDF-1.5 specification allows collecting non-stream objects as "compressed objects" into "object stream" objects (/Type/ObjStm, see the PDF reference manual, 5th ed., §3.4.6). At the end of the PDF file, an /XRef cross-reference stream is then written out instead of the object table. This can result in a considerably smaller PDF file, particularly if lots of annotations and links are used.

The writing of compressed objects is enabled by setting \pdfobjcompresslevel to a value between 1 and 3; it's disabled if 0 (default). Object compression also requires \pdfminorversion ≥ 5 (or \pdfmajorversion ≥ 2), else a warning is given and the compression is disabled. The value of \pdfobjcompresslevel is clipped to the range 0..3, but using values outside this range is not recommended (for future extension).

Values for \pdfobjcompresslevel have the following effects:

- When set to 0, no object streams are generated at all.

- When set to 1, all non-stream objects are compressed with the exception of any objects coming with embedded PDF files ("paranoid" mode, to avoid yet unknown problems), and also the `/Info` dictionary is not compressed for clear-text legibility.

- When set to 2, also all non-stream objects coming with embedded PDF files are compressed, but the `/Info` dictionary is still not compressed.

- Finally, when set to 3, all non-stream objects are compressed, including the `/Info` dictionary (this means that the `/Info` can't be read as clear text any more). If object streams are to be used, currently `\pdfobjcompresslevel=2` is recommended, and is so specified in some distributions, including TeX Live 2010 and later.

Compatibility caveats: PDF files generated with object streams enabled can't be read with (sufficiently old) PDF viewers that don't understand PDF-1.5 files. For widest distribution and unknown audience, don't activate object stream writing. The PDF-1.5 standard describes also a hybrid object compression mode that gives some backward compatibility, but this is currently not implemented, as PDF-1.5 was rather quickly adopted by modern PDF viewers. Also not implemented is the optional `/Extends` key.

4.1.5 \pdfdecimaldigits

`\pdfdecimaldigits`   (integer)

This integer parameter specifies the numeric accuracy of real coordinates as written to the PDF file. It gives the maximal number of decimal digits after the decimal point. Valid values are in the range 0..4. A higher value means more precise output, but also results in a larger file size and more time to display or print. In most cases the optimal value is 2. This parameter does not influence the precision of numbers used in raw PDF code, like that used in `\pdfliteral` and annotation action specifications; also multiplication items (e.g., scaling factors) are not affected and are always output with best precision. This parameter is read when pdfTeX writes a real number to the PDF output.

When including huge MetaPost images using `supp-pdf.tex`, one can limit the accuracy to two digits with `\twodigitMPoutput`.

4.1.6 \pdfhorigin

`\pdfhorigin`   (dimen)

This parameter can be used to set the horizontal offset of the output box from the top left corner of the page. A value of 1 inch corresponds to the normal TeX offset. This parameter is read when pdfTeX starts shipping out a page to the PDF output.

For all normal purposes, this parameter should always be kept at 1 true inch. If you want to shift text on the page, use TeX's own `\hoffset` primitive. To avoid surprises, after global magnification has been changed by the `\mag` primitive, the `\pdfhorigin` parameter should still be 1 true inch, e.g., by setting `\pdfhorigin=1 true in` after the `\mag` setting. Or, you can preadjust the `\pdfhorigin` value before typing `\mag`, so that its value after the `\mag` command ends up at 1 true inch again.

4.1.7 \pdfvorigin

`\pdfvorigin`   (dimen)

This parameter is the vertical companion of `\pdfhorigin`, and the notes above regarding `\mag` and true dimensions apply. Also keep in mind that the TeX coordinate system starts in the top left corner (downward), while PDF coordinates start at the bottom left corner (upward).

4.1.8 \pdfpagewidth

\pdfpagewidth (dimen)

This dimension parameter specifies the page width of the PDF output (the screen, the paper, etc.). pdfTeX reads this parameter when it starts shipping out a page. If magnification has been changed by the \mag primitive, check that this parameter reflects the desired true page width. When part of the page falls off the paper or screen, it's quite possible that this parameter is set wrong.

If the value is set to zero, the page width is calculated as

$$width_{\text{box being shipped out}} + 2 \times (\verb|\horigin| + \verb|\hoffset|).$$

It is not wise to rely on this default calculation, since box widths may vary unexpectedly.

4.1.9 \pdfpageheight

\pdfpageheight (dimen)

Similar to the previous item, this dimension parameter specifies the page height of the PDF output; the notes above apply. If set to zero, the page height will be calculated analogously to the above.

4.2 Document info and catalog

4.2.1 \pdfomitinfodict

\pdfomitinfodict (integer)

If nonzero, omit the /Info dictionary completely, as required by the PDF A-4 standard. The primitive was introduced in pdfTeX 1.40.25.

4.2.2 \pdfinfo

\pdfinfo ⟨general text⟩

This primitive allows the user to specify information for the document info dictionary. Provided information can be extracted from the output PDF by, for instance, the pdfinfo program.

The ⟨general text⟩ is a collection of key-value-pairs. The key names are preceded by a /, and the values, being strings, are given between parentheses. All keys, and the primitive itself, are optional. Possible keys are:
/Title,
/Author,
/Subject,
/Keywords,
/Producer (defaults to pdfTeX-1.40.25),
/Creator (defaults to TeX),
/CreationDate (defaults to current date and time, with time zone),
/ModDate (same default),
/Trapped (defaults to /False,
/PTEX.Fullbanner (defaults to the \pdftexbanner string, q.v.).

/CreationDate and /ModDate are expressed in the form D:YYYYMMDDhhmmssTZ, where 'D:' is a constant string prefix, YYYY is the year, MM is the month, DD is the day, hh is the hour, mm is the minutes, ss is the seconds, and TZ is an optional string denoting the time zone, Z for universal time. For example, this is the Unix epoch, the beginning of 1970-01-01 UTC, in this format: D:19700101000000Z. If the time zone is not UTC, it is given as +HH'mm' or -HH'mm', indicating an

offset of the given hours and minutes from UTC, with the given either later (+) or earlier (-) than UTC. For more details, see the PDF reference manual, §7.9.4.

Multiple appearances of \pdfinfo are concatenated. Usually if a key is given more than once, the first appearance will be used, but this is viewer-dependent. Except for standard TeX macro expansion, pdfTeX does not perform any further operations in the ⟨general text⟩ provided by the user.

Here is an example of using \pdfinfo to include the information not supplied by pdfTeX:

```
\pdfinfo {
    /Title        (example.pdf)
    /Author       (Tom and Jerry)
    /Subject      (Example)
    /Keywords     (mouse, cat)
}
```

4.2.3 \pdfinfoomitdate

\pdfinfoomitdate   (integer)

If nonzero, omit the /CreationDate and /ModDate keys from the document info dictionary described above. This can be useful in making reproducible PDFs. The primitive was introduced in pdfTeX 1.40.17.

4.2.4 \pdfsuppressptexinfo

\pdfsuppressptexinfo   (integer)

This value is treated as a bit mask, specifying which PTEX.* keys to omit from the output:

value	suppresses
1	PTEX.Fullbanner
2	PTEX.FileName
4	PTEX.PageNumber
8	PTEX.InfoDict

Thus, the value -1, setting all the bits, suppresses everything.

PTEX.Fullbanner is included by default in the general document info dictionary, as mentioned under \pdfinfo above. The other PTEX.* keys are included when a PDF is included in the document (and not otherwise), as described in section 6.

This conditional suppression can be useful in making reproducible PDFs. The primitive was introduced in pdfTeX 1.40.17.

4.2.5 \pdfcatalog

\pdfcatalog ⟨general text⟩ [openaction ⟨action spec⟩]

Similar to the document info section is the document catalog, where possible keys are /URI, which specifies the base url of the document, and /PageMode, which determines how the PDF viewer displays the document on startup. The possibilities for the latter are given in this table:

value	meaning
/UseNone	neither outline nor thumbnails visible
/UseOutlines	outline visible
/UseThumbs	thumbnails visible
/FullScreen	full-screen mode

The default /PageMode setting is /UseNone. In full-screen mode, there is no menu bar, window controls, nor any other window present.

After the ⟨general text⟩ of the catalog, a construct `openaction` ⟨action spec⟩ can be given, where `openaction` is a pdfTEX keyword, and ⟨action spec⟩ specifies the action to be taken when opening the document. This ⟨action spec⟩ is the same as for internal links; see section 4.11.

Several settings can be made in one `\pdfcatalog` call, for example:

```
\pdfcatalog {
  /PageMode /FullScreen
} openaction goto page 2 {/Fit}
```

4.2.6 \pdfcreationdate

`\pdfcreationdate` (expandable)

Expands to the date string pdfTEX uses in the info dictionary of the document, e.g., for this file: `D:20240107094633-08'00'`. The primitive was introduced in pdfTEX 1.30.0.

4.2.7 \pdfnames

`\pdfnames` ⟨general text⟩

This primitive inserts the ⟨general text⟩ to the `/Names` array. The text must conform to the specifications as laid down in the PDF reference manual, or the document may be invalid.

4.2.8 \pdftrailer

`\pdftrailer` ⟨general text⟩

This command puts its argument text verbatim into the file trailer dictionary. Example: `\pdftrailer{/mytrlrkey /mytrlrval}`. The primitive was introduced in pdfTEX 1.11a.

4.2.9 \pdftrailerid

`\pdftrailerid` ⟨general text⟩

Use the ⟨general text⟩ to seed the `/ID` value in the trailer, instead of the default combination of the input file name and starting time. If the argument is empty, the `/ID` is omitted entirely. Example: `\pdftrailerid{}`. This can be useful in making reproducible PDFs. The primitive was introduced in pdfTEX 1.40.17.

4.3 Fonts

4.3.1 \pdfadjustspacing

`\pdfadjustspacing` (integer)

This primitive controls whether font expansion happens (this operation is described in detail at `\pdffontexpand`). By default, `\pdfadjustspacing` is set to 0; then font expansion is disabled, so that the pdfTEX output is identical to that from the original TEX engine.

Font expansion can be activated in two modes. When `\pdfadjustspacing` is set to 1, font expansion is applied *after* TEX's normal paragraph breaking routines have broken the paragraph into lines. In this case, line breaks are identical to standard TEX behavior.

When set to 2, the width changes that are the result of stretching and shrinking are taken into account *while* the paragraph is broken into lines. In this case, line breaks are likely to be different from those of standard TEX. Paragraphs may well become longer or shorter.

Both alternatives require a collection of TFM files that are related to the ⟨stretch⟩ and ⟨shrink⟩ settings for the `\pdffontexpand` primitive, unless this is given with the `autoexpand` option.

4.3.2 \pdffontexpand

\pdffontexpand ⟨font⟩ ⟨stretch⟩ ⟨shrink⟩ ⟨step⟩ [`autoexpand`]

This extension to TEX's font definitions controls a major pdfTEX feature called *font expansion.* To enable font expansion, \pdfadjustspacing must be set to a value greater than zero. We describe the basic process with an example:

```
\font\somefont=sometfm at 10pt
\pdffontexpand\somefont 30 20 10 autoexpand
\pdfadjustspacing=2
```

The `30 20 10` means this: "hey TEX, when line breaking is going badly, you may stretch the glyphs from this font as much as 3% or shrink them as much as 2%." For practical reasons pdfTEX uses discrete expansion steps, in this example, 1%.

Roughly speaking, the idea is as follows: When TEX cannot break a line in the appropriate way, the unbreakable parts of the last word may stick into the margin. When pdfTEX sees this, it will try to scale (shrink) the glyphs in that line using fixed steps, until the line fits. When lines are too spacey, the opposite happens: pdfTEX starts scaling (stretching) the glyphs until the white space gaps are acceptable. This glyph stretching and shrinking is called *font expansion.*

There are two different modes for font expansion, depending on whether `autoexpand` is specified:

1. If the `autoexpand` keyword is given—this is recommended mode—only a single map entry is needed for all expanded font versions, using the name of the unexpanded TFM file (*tfmname* in section 3.1.1). No expanded *tfmname* versions need be mentioned (indeed, they are ignored), as pdfTEX generates expanded instances of the unexpanded TFM data structures and keeps them in its memory. As of pdfTEX 1.40.0, the `autoexpand` work is done within the page stream by modification of the text matrix (PDF operator "Tm"), and not at the font file level, giving the advantage that it now works not only with Type 1 but also with TrueType and OpenType fonts (and even without embedding a font file; but that's not recommended). In this mode pdfTEX requires only unexpanded font files.

2. Second, if the `autoexpand` keyword is not given, setting up font expansion requires considerably more work, as there must be map entries for TFM files in all required expansion values. The expanded *tfmname* variants are constructed by adding the font expansion value to the *tfmname* of the base font, e.g., there must be a map entry with *tfmname* `sometfm+10` for 1% stretch or `sometfm-15` for 1.5% shrink. This also means that for each expanded font variant a TFM file with properly expanded metrics must exist. In addition to the TFM file, it is necessary to provide, for each expansion value, an individually crafted font file with the expanded glyphs. Thus, this allows the absolute best possible output, controlling the glyphs for every expanded variant of the font. It is a rare document indeed for this to be worth the trouble.

   One technical drawback of non-`autoexpand` mode is that all needed individual font files need to be embedded in the PDF output for each expanded font, leading to significantly larger PDF files than in `autoexpand` mode.

   Another caveat for non-`autoexpand` mode: when \pdffontexpand is executed, pdfTEX immediately loads two fonts, at the maximum stretch and shrink; in our example, `sometfm+` `30` and `sometfm-20`. (If they aren't available, `mktextfm` may be uselessly called, and then an error message issued.) This happens even if those fonts never end up being used, which is arguably undesirable, but hard to change. It is not a problem when using `autoexpand`.

The font expansion mechanism is inspired by an optimization first introduced by Prof. Hermann Zapf, which in itself goes back to optimizations used in the early days of typesetting: use different glyphs to optimize the grayness of a page. So, there are many, slightly different a's, e's, etc. For practical reasons pdfTeX does not use such huge glyph collections; it uses horizontal scaling instead. This is sub-optimal, and possibly offensive to the design, though the expansions are tiny. It is up to the user to decide whether such slightly remastered fonts are acceptable. As an example, this document is typeset with font expansion and margin kerning activated (via the `microtype` LaTeX package).

4.3.3 \efcode

\efcode ⟨font⟩ ⟨8-bit number⟩   (integer)

We haven't yet told the whole story. One can imagine that some glyphs are visually more sensitive to stretching or shrinking than others. Then the `\efcode` primitive can be used to influence the expandability of individual glyphs within a given font, as a factor applied to the expansion setting from the `\pdffontexpand` primitive. The syntax is similar to `\sfcode` (but with the ⟨font⟩ required), and it defaults to 1000, meaning 100% expandability. The given integer value is clipped to the range 0..1000, corresponding to a usable expandability range of 0..100%. Examples:

```
\efcode\somefont`A=800
\efcode\somefont`O=0
```

Here the A in `\somefont` may only stretch or shrink by up to 80% of the current expansion value for that font, and expansion for the O is disabled. The actual expansion is still bound to the steps as defined by `\pdffontexpand` primitive.

Changes to this table are global, i.e., ignore TeX's usual grouping, and apply only to the given ⟨font⟩.

4.3.4 \pdfprotrudechars

\pdfprotrudechars   (integer)

Yet another way of optimizing paragraph breaking is to let certain characters move into the margin ('character protrusion'). Character protrusion is disabled when `\pdfprotrudechars=0` or negative.

When `\pdfprotrudechars=1`, the glyphs qualified as such will make this move when applicable, without changing the line-breaking. When `\pdfprotrudechars=2` (or greater), character protrusion will be taken into account while considering breakpoints, so line-breaking might be changed. This qualification and the amount of shift are set by the primitives `\rpcode` and `\lpcode`.

If you want to protrude an item other than a character (e.g., an `\hbox`), you can do so by padding the item with an invisible zero-width character for which protrusion is activated.

4.3.5 \rpcode, \lpcode

\rpcode ⟨font⟩ ⟨8-bit number⟩   (integer)
\lpcode ⟨font⟩ ⟨8-bit number⟩   (integer)

The amount that a character from a given font may shift into the right margin ('character protrusion') is set by the primitive `\rpcode`. The protrusion distance is the integer value given to `\rpcode`, multiplied by 0.001em from the current font. The given integer value is clipped to the range −1000..1000, corresponding to a usable protrusion range of −1em..1em. `\lpcode` is exactly analogous to `\rpcode`, but affects the amount by which characters may protrude into the left margin.

Examples:

```
\rpcode\somefont',=200
\rpcode\somefont'-=150
```

Here the comma may shift by 0.2em into the margin and the hyphen by 0.15em. All these small bits and pieces will help pdfTEX to give you better paragraphs (use `\rpcode` judiciously; don't overdo it).

Remark: old versions of pdfTEX use the character width as measure. This was changed to a proportion of the em-width after Hàn Thế Thành finished his master's thesis.

Changes to these tables are global, i.e., ignore TEX's usual grouping, and apply only to the given ⟨font⟩.

4.3.6 \leftmarginkern, \rightmarginkern

`\leftmarginkern` ⟨box number⟩   (expandable)
`\rightmarginkern` ⟨box number⟩   (expandable)

The `\leftmarginkern` ⟨box number⟩ primitive expands to the width of the margin kern at the left side of the horizontal list stored in `\box` ⟨box number⟩. The expansion string includes the unit `pt`. E.g., when the left margin kern of `\box0` amounts to −10pt, `\leftmarginkern0` will expand to `-10pt`. The primitive `\rightmarginkern` works analogously for the right margin. The primitives were introduced in pdfTEX 1.30.0.

These are auxiliary primitives to make character protrusion more versatile. When using the TEX primitives `\unhbox` or `\unhcopy`, the margin kerns at either end of the unpackaged hbox will be removed (e.g., to avoid weird effects if several hboxes are unpackaged behind each other into the same horizontal list). These `\unhbox` or `\unhcopy` commands are often used together with `\vsplit` for dis- and re-assembling of paragraphs, e.g., to add line numbers. Paragraphs treated like this do not show character protrusion by default, as the margin kerns have been removed during the unpacking process.

The `\leftmarginkern` and `\rightmarginkern` primitives allow access to the margin kerns and store them away before unpackaging the hbox. E.g., the following code snippet restores margin kerning of a horizontal list stored in `\box\testline`, resulting in a hbox with the original margin kerning, now inserted by ordinary kerns.

```
\dimen0=\leftmarginkern\testline
\dimen1=\rightmarginkern\testline
\hbox to\hsize{\kern\dimen0\unhcopy\testline\kern\dimen1}
```

4.3.7 \letterspacefont

`\letterspacefont` ⟨control sequence⟩ ⟨font⟩ ⟨integer⟩

This primitive creates an instance of ⟨font⟩ with the widths of all glyphs increased by ⟨integer⟩ thousandths of an em (as defined by ⟨font⟩). The effect is letter spacing, but the glyphs are actually wider, as the sidebearings are increased, so a single glyph will take more space. For instance, the following creates a font `\lsfont` whose glyphs are all 1.2pt larger than those of `\normalfont`:

```
\font\normalfont=myfont at 12pt
\letterspacefont\lsfont\normalfont 100
```

Negative values are allowed for ⟨integer⟩. Letter spacing works natively in PDF mode only, unless special fonts are devised (in our example, a `myfont+100ls` font), as with font expansion.

4.3.8 \pdfcopyfont

\pdfcopyfont ⟨control sequence⟩ ⟨font⟩

This primitive defines ⟨control sequence⟩ as a synonym for ⟨font⟩.

4.3.9 \pdffakespace

\pdffakespace

Always insert a fake interword space in the output (PDF only; this primitive is an error in DVI mode), regardless of the values of \pdfinterwordspaceon and \pdfinterwordspaceoff (q.v.). Example:

```
Text with a fake interword \pdffakespace space.
```

The primitive was introduced in pdfTeX 1.40.15.

4.3.10 \pdffontattr

\pdffontattr ⟨font⟩ ⟨general text⟩

This primitive inserts the ⟨general text⟩ into the /Font dictionary. The text must conform to the specifications as laid down in the PDF reference manual, or the document may be invalid. For examples, see the cmap and CJK packages.

4.3.11 \pdffontname

\pdffontname ⟨font⟩   (expandable)

In PDF files produced by pdfTeX one can recognize a font resource by the prefix /F followed by a number, for instance /F12 or /F54. For a given TeX ⟨font⟩, this primitive expands to the number from the corresponding font resource name. E.g., if /F12 corresponds to some TeX font \foo, then \pdffontname\foo expands to the number 12.

In the current implementation, when \pdfuniqueresname (see below) is set to a positive value, the \pdffontname still returns only the number from the font resource name, without the appended random string.

4.3.12 \pdffontobjnum

\pdffontobjnum ⟨font⟩   (expandable)

This command is similar to \pdffontname, but it returns the PDF object number of the font dictionary instead of the number from the font resource name. E.g., if the font dictionary (/Type/Font) in PDF object 3 corresponds to some TeX font \foo, then \pdffontobjnum\foo gives the number 3.

Use of \pdffontname and \pdffontobjnum allows users full access to all the font resources used in a document.

4.3.13 \pdffontsize

\pdffontsize ⟨font⟩   (expandable)

This primitive expands to the font size of the given font, with unit pt. E.g., when using the plain TeX macro package, the call \pdffontsize\tenrm expands to 10.0pt.

4.3.14 \pdfgentounicode

\pdfgentounicode (integer)

By default, pdfTeX does not include a /ToUnicode resource when including fonts in the output. Such a resource (also called a CMap resource) maps glyph names to Unicode characters in a PDF file. Lacking such a resource, it is the PDF reader which determines how and whether searching in the PDF file works. In practice, searching for basic ASCII characters generally works, but searching for anything beyond those, including ligatures such as 'fi', is likely to fail.

If \pdfgentounicode is set to 1 when the job ends, the /ToUnicode resource will be included in the output, with mappings for Type 1 fonts used—unless \pdfnobuiltintounicode (q.v.) is set for a given font.

The mapping is created as follows: for each glyph in the font, look for its ToUnicode value in a global hash table. By default that global hash table is empty, in which case pdfTeX merely emits a warning. Entries are added to the table using the command \pdfglyphtounicode, described next.

4.3.15 \pdfglyphtounicode

\pdfglyphtounicode ⟨general text⟩ ⟨general text⟩

The first argument is the name of a glyph, the second is a string of Unicode numeric values denoting characters, separated by spaces. For instance:

```
\pdfgentounicode=1
\pdfglyphtounicode{ff}{0066 0066}
```

maps the ff ligature to a pair of f characters (whose code is U+0066, that is, ASCII 0x66).

Once a single \pdfglyphtounicode definition is made, whether it is used or not, another feature comes into play: all glyph names of the form uniXXXX or uXXXX are mapped to the natural U+XXXX. Many fonts use this style of naming.

In addition, the glyphtounicode.tex file (distributed with pdfTeX and other software) contains thousands of such definitions, covering most common glyph names. So, for practical purposes, one would probably want:

```
\input glyphtounicode
\pdfgentounicode=1
```

(LaTeX users could load the cmap package to achieve the same effect.)

By default, these glyph name-to-unicode mappings are global. Thus,

```
\pdfglyphtounicode{abc}{1234}
```

would map the glyph named abc to U+1234 for every font. However, it's possible to make a mapping for a single font using a tfm: prefix:

```
\pdfglyphtounicode{tfm:foo/abc}{5678}
```

means that for the font foo.tfm, only, the glyph abc is mapped to U+5678.

4.3.16 \pdfincludechars

\pdfincludechars ⟨font⟩ ⟨general text⟩ (expandable)

This command causes pdfTeX to treat the characters in ⟨general text⟩ as if they were used with ⟨font⟩, which means that the corresponding glyphs will be embedded into the font resources in the PDF output. Nothing is appended to the list being built.

4.3.17 \pdfinterwordspaceon, \pdfinterwordspaceoff, \pdfspacefont

```
\pdfinterwordspaceon
\pdfinterwordspaceoff
\pdfspacefont ⟨general text⟩
```

The first two commands insert corresponding whatsit nodes which turn on/off generation of faked interword spaces in the PDF output (they cause errors in DVI output). This allows for better reflowing of text on the fly by PDF readers, better extraction of textual content, and is required by PDF/A. It does not affect the normal TeX justification with glue of the typeset output.

This works roughly as follows: with \pdfinterwordspaceon, pdfTeX will guess when an interword space should be inserted, based on movement within some limits in horizontal direction. When found, pdfTeX inserts a true space character into the PDF page description for the viewers, and adjusts the next movement so that the next character will be in the expected position, according to normal TeX behavior.

Where does that "true space character" come from? There are two possibilities.

- If the current font has a real space character, it is used. pdfTeX considers a font to have such a space character if 1) the font has an encoding vector (.enc file) specified in its map entry, and 2) the encoding has a glyph named space (that is, the PostScript name /space) at slot 32. For example, the font texnansi-lmr10 uses the encoding file lm-texnansi.enc, which has such an entry.

- If the current font does not have such a space character (this is the case for most traditional TeX fonts, such as cmr10 and ec-lmr10), pdfTeX will use the space character from a special fallback font named (by default) pdftexspace[.tfm]. pdfTeX automatically defines a map entry for this font which looks like this:

  ```
  \pdfmapline{=pdftexspace PdfTeX-Space <pdftexspace.pfb}
  ```

  The pdftexspace.tfm and pdftexspace.pfb files are expected to be available to pdfTeX just like any other font. (They are distributed with pdfTeX.) The pdftexspace font was constructed by hand; it has a space character that is .333em (and no other characters).

  A different fallback font for the space character can be given via \pdfspacefont{myfont}. This is most likely to be useful for testing and debugging. In this case, pdfTeX assumes that the given font has a real space character at slot 32, and that any necessary corresponding map entry exists. For example:

  ```
  \pdfspacefont{texnansi-lmr10} % use space char from this font if
                                % current font has no space char
  ```

History: Before pdfTeX version 1.40.25, no check was made for a space character in the current font, the fallback font was named dummy-space, and its space character was tiny, 0.001em. It turned out that PDF viewers were unhappy with this tiny space, especially with tagged PDF.

Example of usage (see also the fake-interword-space.tex test file):

```
Text with no interword spaces.

\pdfglyphtounicode{space}{0020}
\pdfinterwordspaceon
```

```
Switch to text with faked interword spaces.

\pdfinterwordspaceoff

Back to text with no interword spaces.
```

The primitives were introduced in pdfTEX 1.40.15, 1.40.25.

4.3.18 \pdfmapfile

\pdfmapfile ⟨map filename⟩

This primitive is used for reading a font map file consisting of one or more font map lines. The name of the map file is given in the ⟨map filename⟩ together with an optional leading modifier character, as explained below. If your map file isn't in the current directory or a standard system directory, you will need to set the `TEXFONTMAPS` variable (in Web2C) or give an explicit path so that it will be found.

If no \pdfmapfile primitive is given, the default map file `pdftex.map` will be read by pdfTEX. Normally there is no need for a pdfTEX user to bother about the \pdfmapfile or \pdfmapline primitives, as the main TEX distributions provide helper tools to automatically assemble the default `pdftex.map`. (In TEX Live, these tools are `updmap` and `updmap-sys`.)

There is a companion primitive \pdfmapline that allows scanning single map lines; its map line argument has the same syntax as the map lines from a map file. Both primitives can be used concurrently. The \pdfmapfile primitive is fast for reading external bulk font map information (many map lines collected in a map file), whereas the \pdfmapline allows putting the font map information for individual TEX fonts directly in the TEX source or a style file. With both primitives, the map line information is scanned by pdfTEX identically. In the most common case, the data are put into a fresh internal map entry data structure, which is then consulted when a font is used.

When a \pdfmapfile or \pdfmapline primitive is executed by pdfTEX, the argument (map file or map line) will be processed immediately, and the internal map entry database updated. The operation mode of the \pdfmapfile and \pdfmapline primitives is selected by an optional modifier character, one of +, =, -, in front of the *tfmname* field. This modifier defines how the individual map lines are going to be handled, and how a collision between an already registered map entry and a newer one is resolved: either by ignoring a later entry, or replacing or deleting an existing entry. In any case, map entries of fonts already in use are kept untouched. Here are two examples:

```
\pdfmapfile{+myfont.map}
\pdfmapline{+ptmri8r Times-Italic <8r.enc <ptmri8a.pfb}
```

When no modifier character is given (e.g., \pdfmapfile{foo.map} or
\pdfmapline{helv Helvetica}) and there has been no previous call to one of these primitives, then the default map file `pdftex.map` will *not* be read. Apart from this case, the given map file will be processed as with the + modifier: duplicate later map entries within the file are ignored and a warning is issued. Thus, you can block reading of the default map file also with an empty \pdfmapfile{} or \pdfmapline{} early in the TEX file. Since the default map file is typically large, if you don't need it, these command variants might considerably speed up pdfTEX startup.

If a modifier is given, before reading the items given as arguments to the present \pdfmapfile or \pdfmapline, the default map file will be read first—if this hasn't already been done or been prevented by the above blocking cases. The meaning of the modifiers:

- \pdfmapfile{+foo.map} reads the file `foo.map`; duplicate later map entries within the file are ignored and a warning is issued.

- `\pdfmapfile{=foo.map}` reads the file `foo.map`; matching map entries in the database are replaced by new entries from `foo.map`, if they haven't already been used.

- `\pdfmapfile{-foo.map}` reads the file `foo.map`; matching map entries are deleted from the database, if they haven't already been used.

In short, if you want to add support for a new font through an additional font map file while keeping all the existing mappings, use `\pdfmapfile{+myfont.map}` or `\pdfmapfile{=myfont.map}`.

If you want to use a base map file name other than `pdftex.map`, or change its processing options through a pdfTEX format, you can do this by appending the `\pdfmapfile` command to the `\everyjob` token list for the `-ini` run, as in:

```
\everyjob=\expandafter{\the\everyjob\pdfmapfile{+myspecial.map}}
\dump
```

This would always read the file `myspecial.map` after the default `pdftex.map` file.

4.3.19   \pdfmapline

`\pdfmapline` ⟨map spec⟩

Similar to `\pdfmapfile`, but here you give a single map line (exactly like the ones in map files) as an argument. The optional modifiers (+-=) have the same effect as with `\pdfmapfile`; see also the description above. Example:

```
\pdfmapline{+ptmri8r Times-Italic <8r.enc <ptmri8a.pfb}
```

This primitive, especially the `\pdfmapline{=...}` form, is useful for temporary quick checks of a new font map entry during development, before finally putting it into a map file.

As explained above, `\pdfmapline{}`, like `\pdfmapfile{}`, blocks reading of the default map file, if it comes early enough in the TEX input. The primitive was introduced in pdfTEX 1.20a.

4.3.20   \pdfmovechars

`\pdfmovechars`   (integer)

Since pdfTEX version 1.30.0 the primitive `\pdfmovechars` is obsolete, and its use merely leads to a warning. (This primitive specified whether pdfTEX should try to move characters in range 0..31 to higher slots; its sole purpose was to remedy certain bugs of early PDF viewers.)

4.3.21   \pdfnobuiltintounicode

`\pdfnobuiltintounicode` ⟨font⟩

The purpose of this command is to prevent pdfTEX from generating the `ToUnicode`/CMap resource for the given font when `\pdfgentounicode=1`, most likely because the CMap resource is already generated by some other method. For instance, the LATEX `cmap` package uses `\pdffontattr` to generate CMap resources.

Minimal example:

```
\font\f=cmb10
\pdfnobuiltintounicode\f
\f No unicode mappings for this output.
```

The primitive was introduced in pdfTEX 1.40.11.

4.3.22 \pdfnoligatures

\pdfnoligatures ⟨font⟩

This disables all ligatures in the loaded font ⟨font⟩. The primitive was introduced in pdfTEX 1.30.0.

4.3.23 \pdfomitcharset

\pdfomitcharset (integer)

If this primitive parameter is zero (the default), the /CharSet entry is included as usual for fonts in the PDF output; if it is set to 1, then /CharSet is omitted. Other values may have other meanings in the future, so do not rely on them.

Explanation: This parameter was created because the PDF/A-1 standard requires /CharSet, whereas PDF/A-2 and PDF/A-3 allow it to be omitted but have extraordinary requirements, which pdfTEX does not currently meet, if it is included.The primitive was introduced in pdfTEX 1.40.20.

4.3.24 \pdfpkmode

\pdfpkmode (tokens)

The \pdfpkmode is a token register that sets the METAFONT mode for pixel font generation. The contents of this register is dumped into the format, so one can (optionally) preset it. The primitive was introduced in pdfTEX 1.30.0.

4.3.25 \pdfpkresolution

\pdfpkresolution (integer)

This integer parameter specifies the default resolution of embedded PK fonts and is read when pdfTEX embeds a PK font during finishing the PDF output. As bitmap fonts may be rendered poorly, and in any case cannot be arbitrarily magnified, it is best to use outline fonts if possible.

4.3.26 \pdfsuppresswarningdupmap

\pdfsuppresswarningdupmap (integer)

Ordinarily, pdfTEX gives a warning when a duplicate map entry for a given font is read, whatever the mechanism. However, sometimes it is useful to include map information within the document, using \pdfmapfile or \pdfmapline, even for fonts that happen to be installed on the system. Then seeing the warnings on every run is just noise; it can be suppressed by setting this parameter to a positive number. The primitive was introduced in pdfTEX 1.40.13.

4.3.27 \pdftracingfonts

\pdftracingfonts (integer)

This integer parameter specifies the level of verbosity for the information about expanded fonts given in the log, e.g., when \tracingoutput=1. If \pdftracingfonts=0, which is the default, the log shows the actual nonzero signed expansion value for each expanded letter within brackets, as in:

```
...\xivtt (+20) t
```

If \pdftracingfonts=1, the name of the TFM file is also listed, together with the font size:

```
...\xivtt (cmtt10+20@14.0pt) t
```

Setting \pdftracingfonts to a value other than 0 or 1 is not recommended, to allow for future extensions. The primitive was introduced in pdfTEX 1.30.0.

4.3.28 \pdfuniqueresname

\pdfuniqueresname (integer)

When this primitive is assigned a positive number, PDF resource names will be made reasonably unique by appending a random string consisting of six ASCII characters.

4.3.29 \tagcode

\tagcode ⟨font⟩ ⟨8-bit number⟩ (integer)

This primitive accesses a character's `char_tag` info. It is meant to delete `lig_tag` (the character's ligature/kerning program), `list_tag` (which indicates that the character belongs to a chain of ascending sizes) and/or `ext_tag` (which indicates that the character is extensible), with the following options: assigning `-7` or less clears all tags, `-6` clears `ext_tag` and `list_tag`, `-5` clears `ext_tag` and `lig_tag`, `-4` clears `ext_tag`, `-3` clears `list_tag` and `lig_tag`, `-2` clears `list_tag`, `-1` clears `lig_tag`, and `0` or larger does nothing. Changes are irreversible and global.

Conversely, when queried, the primitive returns `0` if the tag's value is `no_tag`, `1` if `lig_tag` is set, `2` if `list_tag` is set or `4` (not `3`) if `ext_tag` is set.

4.4 Spacing

Controlling spacing before and after characters was introduced in version 1.30, mostly to handle punctuation rules in different languages. The \...code tables here, like those in the previous section, operate globally, i.e., ignore TeX's usual grouping, and apply only to the given ⟨font⟩, not other instances of the underlying font.

4.4.1 \pdfadjustinterwordglue

\pdfadjustinterwordglue (integer)

If positive, adjustment of interword glue is enabled and controlled by the following three primitives.

4.4.2 \knbscode

\knbscode ⟨font⟩ ⟨8-bit number⟩ (integer)

The amount of space, in thousandths of an em, added to the natural width of the glue following a character (the name stands for "kern before space", although technically it is looking at glue items, not kern items). This amounts is clipped to the range $-1000..1000$. For instance, in the following example, glue after periods in the current font will be increased by .2em.

```
\pdfadjustinterwordglue=1
\knsbcode\font'\.=200
```

4.4.3 \stbscode

\stbscode ⟨font⟩ ⟨8-bit number⟩ (integer)

This works like \knbscode, but applies to the stretch component of the following glue.

4.4.4 \shbscode

\shbscode ⟨font⟩ ⟨8-bit number⟩ (integer)

Like \stbscode, but for the shrink component.

4.4.5 \pdfprependkern

`\pdfprependkern` (integer)

If positive, automatic insertion of kerns before characters is enabled.

4.4.6 \knbccode

`\knbccode` ⟨font⟩ ⟨8-bit number⟩ (integer)

The width of the kern, in thousandths of an em, inserted before a character. It is clipped to the range $-1000..1000$. For instance, with the following code, a .15em-kern will be inserted before all question marks in the current font (possibly useful for e.g., French punctuation):

```
\pdfprependkern=1
\knbccode\font'\?=150
```

4.4.7 \pdfappendkern

`\pdfappendkern` (integer)

Same as `\pdfprependkern`, but for kerns inserted after characters.

4.4.8 \knaccode

`\knaccode` ⟨font⟩ ⟨8-bit number⟩ (integer)

Same as `\knbccode`, except the kern is inserted after the character. Such a kern is required for instance after a left guillemet in French punctuation.

4.5 Vertical adjustments

4.5.1 \pdfignoreddimen

`\pdfignoreddimen` (dimen)

This specifies the dimension value which must be assigned to the following four primitives so they are ignored. Default is `-1000pt`, and it should be modified with care since it also influences when a previous paragraph's depth is ignored (for instance, the plain TeX macro `\nointerlineskip` should be modified accordingly).

4.5.2 \pdffirstlineheight, \pdflastlinedepth

`\pdffirstlineheight` (dimen)
`\pdflastlinedepth` (dimen)

These parameters specify the height of the first, resp. depth of the last, line of a paragraph, regardless of its content. They are read when the paragraph builder is called, and ignored when set to `\pdfignoreddimen`. By default, they are set to `-1000pt`, so they are ignored as long as the value of `\pdfignoreddimen` is not changed.

4.5.3 \pdfeachlineheight, \pdfeachlinedepth

`\pdfeachlineheight` (dimen)
`\pdfeachlinedepth` (dimen)

`\pdfeachlineheight` is similar to `\pdffirstlineheight`, but for all lines of a paragraph, including the first one, unless `\pdffirstlineheight` is specified.

`\pdfeachlinedepth` is the same, but for the depth.

4.6 PDF objects

4.6.1 \pdfobj

\pdfobj ⟨object type spec⟩ (h, v, m)

This command creates a raw PDF object that is written to the PDF file as `1 0 obj ... endobj`.

When ⟨object type spec⟩ is not given, pdfTeX no longer creates a dictionary object with contents ⟨general text⟩, as it did in the past.

When ⟨object type spec⟩ is given as ⟨attr spec⟩ `stream`, the object will be created as a stream with contents ⟨general text⟩ and additional attributes in ⟨attr spec⟩.

When ⟨object type spec⟩ is given as ⟨attr spec⟩ `file`, then the ⟨general text⟩ will be treated as a file name and its contents will be copied into the stream contents.

When ⟨object type spec⟩ is given as `reserveobjnum`, just a new object number is reserved. The number of the reserved object is accessible via \pdflastobj. The object can later be filled with contents by \pdfobj useobjnum ⟨number⟩ {⟨balanced text⟩}

. But the reserved object number can be used by other objects before it is defined, which provides a forward-referencing mechanism.

The object is kept in memory and will be written to the PDF output only when its number is referred to by \pdfrefobj or when \pdfobj is preceded by \immediate. Nothing is appended to the list being built. The number of the most recently created object is accessible via \pdflastobj.

4.6.2 \pdflastobj

\pdflastobj   (read-only integer)

This command returns the object number of the last object created by \pdfobj.

4.6.3 \pdfrefobj

\pdfrefobj ⟨object number⟩ (h, v, m)

This command appends a whatsit node to the list being built. When the whatsit node is searched at shipout time, pdfTeX will write the object ⟨object number⟩ to the PDF output if it has not been written yet.

4.6.4 \pdfretval

\pdfretval   (read-only integer)

Set to −1 if \pdfobj ignores an invalid object number. Perhaps this will be used to store the error status of other primitives in the future.

4.7 Page and pages objects

4.7.1 \pdfpagesattr

\pdfpagesattr   (tokens)

pdfTeX expands this token list when it finishes the PDF output and adds the resulting character stream to the root `Pages` object. When defined, these are applied to all pages in the document. Some examples of attributes are `/CropBox`, the rectangle specifying the region of the page being displayed and printed, and `/Rotate`, the number of degrees (in multiples of 90) the page should be rotated clockwise when it is displayed or printed.

```
\pdfpagesattr
  { /Rotate 90                 % rotate all pages by 90 degrees
    /CropBox [0 0 612 792] }   % the crop size of all pages (in bp)
```

4.7.2 \pdfpageattr

\pdfpageattr (tokens)

This is similar to \pdfpagesattr, but has priority over it. It can be used to override any attribute given by \pdfpagesattr for an individual page. The token list is expanded when pdfTeX ships out a page. The contents are added to the attributes of the current page.

If the \pdfpageattr value contains the string /MediaBox, then pdfTeX omits outputting its own /MediaBox value (which is [0 0 ⟨page_width⟩ ⟨page_height⟩]). (This behavior was introduced in version 1.40.18.)

4.7.3 \pdfomitprocset

\pdfomitprocset (integer)

If this parameter is zero (the default), the /ProcSet array is included if \pdfmajorversion is 1, and omitted if \pdfmajorversion ≥ 2. If this parameter is > 0, /ProcSet is always omitted; if it is < 0, /ProcSet is always included. For information about what /ProcSet is, see the PDF reference manual or other documentation.

\ProcSet was considered obsolete as of PDF 1.4, but conforming writers should continue to output it. It was formally deprecated in PDF 2.0.The primitive was introduced in pdfTeX 1.40.25.

4.7.4 \pdfpageref

\pdfpageref ⟨page number⟩ (expandable)

This primitive expands to the number of the page object that contains the dictionary for page ⟨page number⟩. If the page ⟨page number⟩ does not exist, a warning will be issued, a fresh unused PDF object will be generated, and \pdfpageref will expand to that object number.

E.g., if the dictionary for page 5 of the TeX document is contained in PDF object no. 18, \pdfpageref5 expands to the number 18.

4.7.5 \pdfpageresources

\pdfpageresources (tokens)

These tokens are added to the resource dictionary for all pages, before the font, XObject, and ProcSet resources. For example:

\pdfpageresources{/MyPageResourceAttribute /MyValue}

4.8   Form XObjects

The next three primitives support a PDF feature called "object reuse" in pdfTeX. The idea is first to create a 'form XObject'. The content of this object corresponds to the content of a TeX box; it can contain pictures and references to other form XObjects as well. After creation, the form XObject can be used multiple times by simply referring to its object number. This feature can be useful for large documents with many similar elements, to reduce the duplication of identical objects.

These commands behave similarly to \pdfobj, \pdfrefobj and \pdflastobj (described in the previous section), but instead of taking raw PDF code, they handle text typeset by TeX.

4.8.1 \pdfxform

\pdfxform [⟨attr spec⟩] [⟨resources spec⟩] ⟨box number⟩ (h, v, m)

This command creates a form XObject corresponding to the contents of the box ⟨box number⟩. The box can contain other raw objects, form XObjects, or images as well. However, it cannot contain annotations because those are laid out on a separate layer, are positioned absolutely, and have dedicated housekeeping. \pdfxform makes box ⟨box number⟩ void, as \box does.

When ⟨attr spec⟩ is given, the text will be written as additional attribute into the form XObject dictionary. The ⟨resources spec⟩ is similar, but the text will be added to the resources dictionary of the form XObject. The text given by ⟨attr spec⟩ or ⟨resources spec⟩ is written before other entries of the form dictionary and/or the resources dictionary and takes priority over later ones.

4.8.2 \pdfrefxform

\pdfrefxform ⟨object number⟩ (h, v, m)

The form XObject is kept in memory and will be written to the PDF output only when its object number is referred to by \pdfrefxform or when \pdfxform is preceded by \immediate. Nothing is appended to the list being built. The number of the most recently created form XObject is accessible via \pdflastxform.

When issued, \pdfrefxform appends a whatsit node to the list being built. When the whatsit node is searched at shipout time, pdfTEX will write the form ⟨object number⟩ to the PDF output if it is not written yet.

4.8.3 \pdflastxform

\pdflastxform (read-only integer)

The object number of the most recently created form XObject is accessible via \pdflastxform.

As said, this feature can be used for reusing information. This mechanism also plays a role in typesetting fill-in forms. Such widgets sometimes have visuals that show up on user request, but are hidden otherwise.

4.8.4 \pdfxformname

\pdfxformname ⟨object number⟩ (expandable)

In PDF files produced by pdfTEX one can recognize a form XObject by the prefix /Fm followed by a number, for instance /Fm2. For a given form XObject number, this primitive expands to the number in the corresponding form XObject name. E.g., if /Fm2 corresponds to some form XObject with object number 7, the \pdfxformname7 expands to the number 2. The primitive was introduced in pdfTEX 1.30.0.

4.9 Graphics inclusion

PDF provides a mechanism for embedding graphic and textual objects: form XObjects. In pdfTEX this mechanism is accessed by means of \pdfxform, \pdflastxform and \pdfrefxform (described in the previous section). A special kind of XObjects are bitmap graphics and for manipulating them similar commands are provided.

4.9.1 \pdfximage

\pdfximage [⟨rule spec⟩] [⟨attr spec⟩] [⟨page spec⟩] [⟨pdf box spec⟩] [⟨colorspace spec⟩]
      ⟨general text⟩ (h, v, m)

This command creates an image object from the filename given in ⟨general text⟩, i.e., enclosed in braces.

The image is kept in memory when this command is executed. It is written to the PDF output only when either 1) the \pdfximage command is preceded by \immediate, or 2) its number is referred to by \pdfrefximage (q.v.).

The dimensions of the image can be controlled via ⟨rule spec⟩. The default values are zero for depth and "running" for height and width. If all of them are given, the image will be scaled to fit the specified values. If some (but not all) are given, the rest will be set to a value corresponding to the remaining ones so as to make the image size to yield the same proportion of $width : (height+depth)$ as the original image size, where depth is treated as zero. If none are given then the image will have its natural size.

An image inserted at its natural size often has a resolution of \pdfimageresolution (see below) given in dots per inch in the output file, but some images may contain data specifying the image resolution, and in such a case the image will be scaled accordingly. The dimensions of an image can be accessed by putting the \pdfrefximage command in a box and checking the box's dimensions:

\setbox0=\hbox{\pdfximage{somefile.png}\pdfrefximage\pdflastximage}

Now we can use \wd0 and \ht0 to find the natural size of the image as determined by pdfTeX.

The image type is specified by the extension of the given file name: .png stands for PNG image, .jpg (or .jpeg) for JPEG, .jbig2 (preferred, but .jb2 works also) for JBIG2, and .pdf for PDF file. But once pdfTeX has opened the file, it checks the file type first by checking the magic number at the file start, which gets precedence over the file name extension. This gives a certain degree of fault tolerance, if the file name extension is stated wrongly.

Similarly to \pdfxform, the optional text given by ⟨attr spec⟩ will be written as additional attributes of the image before other keys of the image dictionary. When using this, be aware that slightly different types of PDF objects are created while including PNG, JPEG, or JBIG2 bitmaps and PDF images.

While working with PDF or JBIG2 images, ⟨page spec⟩ allows to specify which page of the document is to be included; the ⟨page spec⟩ is irrelevant for the other two image formats.

Starting with pdfTeX 1.11, in the PDF image case, you may also choose which page box of the image is to be treated as a final bounding box. If ⟨pdf box spec⟩ is present, it overrides the default behavior specified by the \pdfpagebox parameter, and is overridden by the (obsolete) \pdfforcepagebox parameter. This option is irrelevant for non-PDF inclusions.

Starting with pdfTeX 1.21, \pdfximage supports a colorspace keyword followed by an object number (user-defined colorspace for the image being included). This feature works for JPEG images only. PNG images use RGB palettes, JBIG2 bi-tonal, and PDF images have self-contained color space information.

4.9.2 \pdfrefximage

\pdfrefximage ⟨object number⟩

\pdfrefximage appends a whatsit node to the list being built. When the whatsit node is searched at shipout time, pdfTeX will write the image with number ⟨object number⟩ to the PDF output if it has not been written yet.

4.9.3 \pdflastximage

\pdflastximage   (read-only integer)

\pdflastximage is the number of the most recently created XObject image.

4.9.4 \pdfximagebbox

\pdfximagebbox ⟨integer⟩ ⟨integer⟩   (expandable)

The dimensions of the bounding box of a PDF image loaded with \pdfximage are stored in a table. This primitive returns those dimensions as follows:

```
\pdfximage{afile.pdf}
\pdfximagebbox\pdflastximage 1 % Returns lower-left x
\pdfximagebbox\pdflastximage 2 % Returns lower-left y
\pdfximagebbox\pdflastximage 3 % Returns upper-right x
\pdfximagebbox\pdflastximage 4 % Returns upper-right y
```

4.9.5 \pdflastximagecolordepth

\pdflastximagecolordepth   (read-only integer)

The color depth (1 for 1-bit images, 2 for 2-bit images, and so on) of the last image accessed with \pdfximage.

4.9.6 \pdflastximagepages

\pdflastximagepages   (read-only integer)

This read-only register returns the highest page number from a file previously accessed via the \pdfximage command. This is useful only for PDF files; it always returns 1 for PNG, JPEG, or JBIG2 files.

4.9.7 \pdfimageresolution

\pdfimageresolution   (integer)

The integer \pdfimageresolution parameter (given in dots per inch, dpi) is a last resort value, used only for bitmap (JPEG, PNG, JBIG2) images, but not for PDFs. The priorities are as follows: Often one image dimension (width or height) is stated explicitly in the TEX file. Then the image is properly scaled so that the aspect ratio is kept. If both image dimensions are given, the image will be stretched accordingly, whereby the aspect ratio might get distorted. If no image dimensions are given in the TEX file, the image size will be calculated from its width and height in pixels, using the x and y resolution values normally contained in the image file. If one of these resolution values is missing or weird (either < 0 or > 65535), the \pdfimageresolution value will be used for both x and y resolution when calculating the image size. And if the \pdfimageresolution is zero, a final fallback resolution of 72dpi is used.

The \pdfimageresolution is read when pdfTEX creates an image via \pdfximage. The given value is clipped to the range 0..65535 (dpi).

Currently this parameter is used particularly for calculating the dimensions of JPEG images with EXIF information (unless at least one dimension is stated explicitly); the resolution values included in EXIF data are ignored.

4.9.8 \pdfpagebox

`\pdfpagebox` (integer)

When PDF files are included, the command `\pdfximage` allows the selection of which PDF page box to use in the optional ⟨image attr spec⟩. If the option isn't present, the page box defaults to the value of `\pdfpagebox` as follows: (1) media box, (2) crop box, (3) bleed box, (4) trim box, and (5) artbox.

4.9.9 \pdfforcepagebox

`\pdfforcepagebox` (integer)

The integer primitive `\pdfforcepagebox` allows globally overriding the choice of the page box used with `\pdfximage`. It takes the same values as `\pdfpagebox`. The command is available starting with pdfTeX 1.30.0, as a shortened synonym of obsolete `\pdfoptionalwaysusepdfpagebox` instruction, but is itself now considered obsolete—a combination of `\pdfpagebox` and ⟨image attr spec⟩ is better.

4.9.10 \pdfinclusionerrorlevel

`\pdfinclusionerrorlevel` (integer)

This controls the behavior of pdfTeX when a PDF file is included which has a newer PDF version than the one specified by `\pdfmajorversion` and `\pdfminorversion`. If `\pdfinclusionerrorlevel` is set to 0 (the default), pdfTeX gives a warning; if 1, pdfTeX raises an error; if negative, no diagnostic at all is given.

It was originally a shortened synonym of `\pdfoptionpdfinclusionerrorlevel`, which is now obsolete. The primitive was introduced in pdfTeX 1.30.0.

4.9.11 \pdfimagehicolor

`\pdfimagehicolor` (integer)

This parameter, when set to 1, enables embedding of PNG images with 16 bit wide color channels at their full color resolution. This embedding mode is defined only from PDF version 1.5 onwards, so the `\pdfimagehicolor` functionality is automatically disabled in pdfTeX if `\pdfminorversion` < 5 and `\pdfmajorversion` $= 1$; in this case, each 16 bit color channel is reduced to a width of 8 bits by stripping the lower 8 bits before embedding. The same stripping happens when `\pdfimagehicolor` is set to 0. If `\pdfmajorversion` $= 1$ and `\pdfminorversion` ≥ 5, or `\pdfmajorversion` ≥ 2, the default value of `\pdfimagehicolor` is 1.

If specified, the parameter must appear before any data is written to the PDF output. The primitive was introduced in pdfTeX 1.30.0.

4.9.12 \pdfimageapplygamma

`\pdfimageapplygamma` (integer)

This primitive, when set to 1, enables gamma correction while embedding PNG images, taking the value of the primitive `\pdfgamma` as well as the gamma value embedded in the PNG image into account. When `\pdfimageapplygamma` is set to 0, no gamma correction is performed.

If specified, the parameter must appear before any data is written to the PDF output. The primitive was introduced in pdfTeX 1.30.0.

4.9.13 \pdfgamma

`\pdfgamma` (integer)

This primitive defines the 'device gamma' for pdfTeX. Values are in promilles (same as `\mag`). The default value of this primitive is 1000, defining a device gamma value of 1.0.

If `\pdfimageapplygamma` is set to 1, then whenever a PNG image is included, pdfTeX applies a gamma correction. This correction is based on the value of the `\pdfgamma` primitive and the 'assumed device gamma' that is derived from the value embedded in the actual image. If no embedded value can be found in the PNG image, then the value of `\pdfimagegamma` is used instead.

If specified, the parameter must appear before any data is written to the PDF output. The primitive was introduced in pdfTeX 1.30.0.

4.9.14 \pdfimagegamma

`\pdfimagegamma` (integer)

This primitive gives a default 'assumed gamma' value for PNG images. Values are in promilles (same as for `\pdfamma`). The default value of this primitive is 2200, implying an assumed gamma value of 2.2.

When pdfTeX is applying gamma corrections, images that do not have an embedded 'assumed gamma' value are assumed to have been created for a device with a gamma of 2.2. Experiments show that this default setting is correct for a large number of images; however, if your images come out too dark, you probably want to set `\pdfimagegamma` to a lower value, like 1000.

If specified, the parameter must appear before any data is written to the PDF output. The primitive was introduced in pdfTeX 1.30.0.

4.9.15 \pdfpxdimen

`\pdfpxdimen` (dimen)

While working with bitmap graphics or typesetting electronic documents, it might be convenient to base dimensions on pixels rather than TeX's standard units like `pt` or `em`. For this purpose, pdfTeX provides an extra unit called `px` that takes the dimension given to the `\pdfpxdimen` primitive. In example, to make the unit `px` corresponding to 96dpi pixel density (then 1px = 72/96bp), one can do the following calculation:

```
\pdfpxdimen=1in % 1 dpi
\divide\pdfpxdimen by 96 % 96 dpi
\hsize=1200px
```

Then `\hsize` amounts to 1200 pixels in 96dpi, which is exactly 903.375pt (but TeX rounds it to 903.36914pt). The default value of `\pdfpxdimen` is 1bp, corresponding to a pixel density of 72dpi. This primitive is completely independent from the `\pdfimageresolution` and `\pdfpkresolution` parameters.

The primitive was introduced in pdfTeX 1.30.0. It used to be an integer register that gave the dimension 1px as number of scaled points, defaulting to 65536 (1px equal to 65536sp = 1pt). Starting with pdfTeX 1.40.0, `\pdfpxdimen` is now a real dimension parameter.

4.9.16 \pdfinclusioncopyfonts

`\pdfinclusioncopyfonts` (integer)

If positive, this parameter tells pdfTeX to include fonts that are embedded in a PDF file loaded with `\pdfximage`, even if those fonts are also available on disk. Bigger files might be created, but

included PDF files are sure to be embedded with the correct fonts; indeed, the fonts on disk might be different from the embedded ones, and glyphs might be missing.

4.9.17 \pdfsuppresswarningpagegroup

`\pdfsuppresswarningpagegroup` (integer)

Ordinarily, pdfTeX gives a warning when more than one included PDF file has a so-called "page group object" (`/Group`), because only one can "win"—that is, be propagated to the page level. Usually the page groups are identical, but when they are not, the result is unpredictable. It would be ideal if pdfTeX in fact detected whether the page groups were the same and only gave the warning in the problematic case; unfortunately, this is not easy (a patch would be welcome). Nevertheless, often one observes that there is no actual problem. Then seeing the warnings on every run is just noise, and can be suppressed by setting this parameter to a positive number. The primitive was introduced in pdfTeX 1.40.15.

4.10 Annotations

PDF 1.4 provides four basic kinds of annotations:

- hyperlinks, general navigation
- text clips (notes)
- movies
- sound fragments

The first type differs from the other three in that there is a designated area involved on which one can click, or when moved over some action occurs. pdfTeX is able to calculate this area, as we will see later. All annotations can be supported using the next two general annotation primitives.

4.10.1 \pdfannot

`\pdfannot` ⟨annot type spec⟩ (h, v, m)

This command appends a whatsit node corresponding to an annotation to the list being built. The dimensions of the annotation can be controlled via the ⟨rule spec⟩. The default values are "running" for all width, height and depth. When an annotation is written out, running dimensions will take the corresponding values from the box containing the whatsit node representing the annotation. The ⟨general text⟩ is inserted as raw PDF code to the contents of annotation. The annotation is written only if the corresponding whatsit node is searched at shipout time.

4.10.2 \pdflastannot

`\pdflastannot` (read-only integer)

This primitive returns the object number of the last annotation created by `\pdfannot`. These two primitives allow users to create any annotation that cannot be created by `\pdfstartlink` (see below).

4.11 Destinations and links

The first type of annotation (hyperlinks and navigation), mentioned above, is implemented by three primitives. The first one is used to define a specific location as being referred to. This location is tied to the page, not the exact location on the page. The main reason for this is that pdfTeX maintains a dedicated list of these annotations, and more when optimized, for the sole purpose of speed.

4.11.1 \pdfdest

\pdfdest ⟨dest spec⟩ (h, v, m)

This primitive appends a whatsit node which establishes a destination for links and bookmark outlines; the link is identified by either a number or a symbolic name, and the way the viewer is to display the page is specified in ⟨dest spec⟩, which must be one of those mentioned in the table below.

keyword	meaning
fit	fit the page in the window
fith	fit the width of the page
fitv	fit the height of the page
fitb	fit the BoundingBox of the page
fitbh	fit the width of BoundingBox of the page
fitbv	fit the height of BoundingBox of the page
fitr ⟨rule spec⟩	fit according to ⟨rule spec⟩
xyz [zoom ⟨number⟩]	goto the current position (see below)
struct [⟨number⟩]	create structure destination (see below)

The xyz keyword can optionally be followed by zoom ⟨integer⟩ to provide a fixed zoom-in. The ⟨integer⟩ is used like the TeX magnification value, i.e., 1000 is the normal page view. When zoom ⟨integer⟩ is given, the zoom factor changes to 0.001 of the ⟨integer⟩ value, otherwise the current zoom factor is kept unchanged.

If struct ⟨number⟩ is used, a structure destination is created instead of a regular destination, referring to the structure element defined in object ⟨number⟩. Structure destinations use a separate namespace and therefore may have the same identifiers as a regular destination.

The destination is written out only if the corresponding whatsit node is searched at shipout time.

4.11.2 \pdfstartlink

\pdfstartlink [⟨rule spec⟩] [⟨attr spec⟩] ⟨action spec⟩ (h, m)

This primitive is used along with \pdfendlink and appends a whatsit node corresponding to the start of a hyperlink. The whatsit node representing the end of the hyperlink is created by \pdfendlink. The dimensions of the link are handled in the similar way as in \pdfannot. Both \pdfstartlink and \pdfendlink must be in the same level of box nesting. A hyperlink with running width can be multi-line or even multi-page, in which case all horizontal boxes with the same nesting level as the boxes containing \pdfstartlink and \pdfendlink will be treated as part of the hyperlink. The hyperlink is written out only if the corresponding whatsit node is searched at shipout time.

Additional attributes, which are explained in great detail in the PDF reference manual, can be given via ⟨attr spec⟩. Typically, the attributes specify the color and thickness of any border around the link. As a simple example, /C [0.9 0 0] /Border [0 0 2] specifies a color (in RGB) of dark red, and a border thickness of 2 points.

While all graphics and text in a PDF document have relative positions, annotations have internally hard-coded absolute positions. Again this is for the sake of speed optimization. The main disadvantage is that these annotations do *not* obey transformations issued by \pdfliteral.

The ⟨action spec⟩ specifies the action that should be performed when the hyperlink is activated, one of (see the syntax rules): user ⟨user-action spec⟩, goto ⟨goto-action spec⟩, thread ⟨thread-action spec⟩.

- A ⟨user-action spec⟩ (user{...}) performs a user-defined action. Examples: a url, as in /S /URI /URI (https://tug.org/); a named action, as in /S /Named /N /NextPage.

- A ⟨goto-action spec⟩ (goto...) performs various goto actions, and is by far the most complex action. Here is a copy of the syntax, for easier reference:

  ⟨goto-action spec⟩ → [⟨goto-action struct spec⟩] ⟨numid⟩ |
  　　　　　[⟨file spec⟩] [⟨goto-action struct spec⟩] ⟨nameid⟩ |
  　　　　　[⟨file spec⟩] [⟨goto-action struct spec⟩] [⟨page spec⟩] ⟨general text⟩ |
  　　　　　⟨file spec⟩ [⟨goto-action struct spec⟩] ⟨nameid⟩ ⟨newwindow spec⟩ |
  　　　　　⟨file spec⟩ [⟨goto-action struct spec⟩] [⟨page spec⟩] ⟨general text⟩ ⟨newwindow spec⟩

  - The ⟨numid⟩ and ⟨nameid⟩ alternatives specify a destination identifier.
  - The ⟨page spec⟩ alternative specifies a page number for the destination. In this case, the zoom factor is given by ⟨general text⟩.
  - The default behavior of ⟨newwindow spec⟩ depends on the browser setting.
  - A destination in another PDF file can be given with ⟨file spec⟩, in which case ⟨newwindow spec⟩ specifies whether the file should be opened in a new window. A ⟨file spec⟩ can be either a (⟨string⟩) or a <<⟨dictionary⟩>>.
  - If a ⟨goto-action spec⟩ contains a ⟨goto-action struct spec⟩, then a structure destination is referenced in addition to the regular destination. The form with ⟨general text⟩ is used if and only if a ⟨file spec⟩ is present; then the ⟨general text⟩ should expand to a literal PDF dictionary describing a structure destination. Otherwise, the ⟨numid⟩ or ⟨nameid⟩ directly after the struct keyword identify a destination which must have been created with \pdfdest struct.

- A ⟨thread-action spec⟩ (thread...) performs article thread reading. The thread identifier, ⟨numid⟩ or ⟨nameid⟩, is similar to the destination identifier. A thread in another PDF file can be referenced by specifying a ⟨file spec⟩.

4.11.3 \pdfendlink

\pdfendlink (h, m)

This primitive ends a link started with \pdfstartlink. All text between \pdfstartlink and \pdfendlink will be treated as part of this link. pdfTeX may break the result across lines (or pages), in which case it will make several links with the same content.

4.11.4 \pdflastlink

\pdflastlink   (read-only integer)

This primitive returns the object number of the last link created by \pdfstartlink (analogous to \pdflastannot). The primitive was introduced in pdfTeX 1.40.0.

4.11.5 \pdflinkmargin

\pdflinkmargin   (dimen)

This dimension parameter specifies the margin of the box representing a hyperlink and is read when a page containing hyperlinks is shipped out.

4.11.6 \pdfdestmargin

\pdfdestmargin   (dimen)

Margin added to the dimensions of the rectangle around the destinations.

4.11.7 \pdfsuppresswarningdupdest

\pdfsuppresswarningdupdest   (integer)

Ordinarily, pdfTeX gives a warning when the same destination is used more than once. However, due to problematic macro packages, sometimes a document may end up producing the warning through no fault of its own, and fixing the macros may not be feasible. Then seeing the warnings on every run is just noise, and can be suppressed by setting this parameter to a positive number. The primitive was introduced in pdfTeX 1.40.13.

4.11.8 \pdfrunninglinkon, \pdfrunninglinkoff

\pdfrunninglinkoff
\pdfrunninglinkon

These commands create corresponding whatsit nodes which turn off/on generation of running links. Their typical usage is to turn off generation of running links in the header or footer of a page. Generation of running links is on when the shipout routine begins.

The generation of running links works roughly like this: pdfTeX keeps a stack of links created by \pdfstartlink, called pdf_link_stack. When writing out an hbox to PDF, pdfTeX checks if the nesting level of the box is the same as the nesting level of the top entry in pdf_link_stack; if so, that box would become a link, too.

The whatsit nodes created by the above primitives turn off/on a flag which controls if a hbox being shipped can become a link, in addition to the previous condition.

Thus, the commands must be inserted before the hbox in question. For example:

```
% (1) good:
\hbox{\pdfrunninglinkoff
  \hbox{text that would become a link otherwise}
}

% (2) bad:
\hbox{\pdfrunninglinkoff text that would become a link otherwise}
% too late; \pdfrunninglinkoff must be inserted before the box
```

4.12 Bookmarks

4.12.1 \pdfoutline

\pdfoutline [⟨attr spec⟩] ⟨action spec⟩ [count ⟨integer⟩] ⟨general text⟩ (h, v, m)

This primitive creates an outline (or bookmark) entry. The first parameter, ⟨attr spec⟩, specifies the action to be taken, and is the same as that allowed for \pdfstartlink.

The `count` ⟨integer⟩ specifies the number of direct subentries under this entry; specify 0 or omit the clause if this entry has no subentries. If the ⟨integer⟩ is negative, then all subentries will be closed and the absolute value of this number specifies the number of subentries.

The ⟨general text⟩ is what will be shown in the outline window. The outline is written to the PDF output immediately.

4.13 Article threads

4.13.1 \pdfthread

\pdfthread [⟨rule spec⟩] [⟨attr spec⟩] ⟨id spec⟩ (h, v, m)

Defines a bead within an article thread. Thread beads with same identifiers (spread across the document) will be joined together.

4.13.2 \pdfstartthread

\pdftstartthread [⟨rule spec⟩] [⟨attr spec⟩] ⟨id spec⟩ (v, m)

This uses the same syntax as \pdfthread, apart that it must be followed by a \pdfendthread. \pdfstartthread and the corresponding \pdfendthread must end up in vboxes with the same nesting level; all vboxes between them will be added into the thread.

In the output routine, if there are other newly created boxes which have the same nesting level as the vbox(es) containing \pdfstartthread and \pdfendthread, they will be also added into the thread, which is probably not what you want. To avoid such undesired behavior, it's often enough to wrap boxes that shouldn't belong to the thread by a box to change their box nesting level.

4.13.3 \pdfendthread

\pdfendthread (v, m)

This ends an article thread started before by \pdfstartthread.

4.13.4 \pdfthreadmargin

\pdfthreadmargin (dimen)

Specifies a margin to be added to the dimensions of a bead within an article thread.

4.14 Literals and specials

4.14.1 \pdfliteral

\pdfliteral [shipout] [direct | page] ⟨general text⟩ (h, v, m)

Analogous to \special in the original TeX, this command inserts raw PDF code into the output, appending a whatsit node to the list being built. This allows support of color and text transformation, among other things.

By default, ⟨general text⟩ is expanded immediately, when the whatsit node is created, as with \special. Starting with pdfTeX 1.40.25, the optional keyword shipout can be used to delay expansion of ⟨general text⟩ until the whatsit node is shipped out, as with non-\immediate \write.

Normally, pdfTeX ends a text section in the PDF output and sets the transformation matrix to the current location on the page before inserting ⟨general text⟩; this can be turned off by giving the optional keyword direct.

Starting with version 1.30.0, pdfTeX supports the keyword page in addition to direct. Both modify the default behavior of \pdfliteral, avoiding translation of the coordinate space before

inserting the literal code. The difference is that the `page` keyword instructs pdfTeX to close a `BT...ET` text block before inserting anything. This means that the literal code inserted refers to the PDF origin (lower-left corner of the page) and can be safely enclosed with `q...Q`. In contrast, using `q...Q` operators inside `\pdfliteral` with the `direct` keyword will produce corrupted PDF output, as the PDF standard doesn't allow doing anything like this within a `BT...ET` block.

4.14.2 \special

`\special {pdf: ⟨text⟩}`

This is equivalent to `\pdfliteral {⟨text⟩}`

.

4.14.3 \special direct

`\special {pdf:direct: ⟨text⟩}`

This is equivalent to `\pdfliteral direct {⟨text⟩}`

.

4.14.4 \special page

`\special {pdf:page: ⟨text⟩}`

This is equivalent to `\pdfliteral page {⟨text⟩}`

.

4.14.5 \special shipout

`\special [shipout] {⟨text⟩}`

Starting with version 1.40.25, pdfTeX extends the `\special` primitive to support the optional keyword `shipout`. This delays expansion of ⟨text⟩ until the page is shipped out, as with non-`\immediate` `\write`. By default, ⟨text⟩ is expanded immediately.

`\special shipout {pdf: ⟨text⟩}`

is equivalent to `\pdfliteral shipout {⟨text⟩}`

.

4.15 Strings

4.15.1 \pdfescapestring

`\pdfescapestring ⟨general text⟩`   (expandable)

Starting with version 1.30.0, pdfTeX provides a mechanism for converting a general text into PDF string. Many characters that may be needed inside such a text (notably parentheses), have a special meaning inside a PDF string object and thus, can't be used literally. The primitive replaces each special PDF character by its literal representation by inserting a backslash before that character. Some characters (e.g., space) are also converted into 3-digit octal number. For example, `\pdfescapestring{Text (1)}` will be expanded to `Text\040\(1\)`. This ensures a literal interpretation of the text by the PDF viewer. The primitive was introduced in pdfTeX 1.30.0.

4.15.2 \pdfescapename

\pdfescapename ⟨general text⟩ (expandable)

Analogous to \pdfescapestring, \pdfescapename replaces each special PDF character inside the general text by its hexadecimal representation preceded by a # character. This ensures the proper interpretation of the text if used as a PDF name object. In example, `Text (1)` will be replaced by `Text#20#281#29`. The primitive was introduced in pdfTeX 1.30.0.

4.15.3 \pdfescapehex

\pdfescapehex ⟨general text⟩ (expandable)

This command converts each character of ⟨general text⟩ into its hexadecimal representation. Each character of the argument becomes a pair of hexadecimal digits. The primitive was introduced in pdfTeX 1.30.0.

4.15.4 \pdfunescapehex

\pdfunescapehex ⟨general text⟩ (expandable)

This command treats each character pair of ⟨general text⟩ as a hexadecimal number and returns the corresponding characters. The primitive was introduced in pdfTeX 1.30.0.

4.15.5 \pdfstrcmp

\pdfstrcmp ⟨general text⟩ ⟨general text⟩ (expandable)

This command compares two strings and expands to 0 if the strings are equal, to -1 if the first string ranks before the second, and to 1 otherwise. The primitive was introduced in pdfTeX 1.30.0.

4.15.6 \pdfmatch

\pdfmatch [icase] [subcount ⟨integer⟩] ⟨general text⟩ ⟨general text⟩ (expandable)

This command implements pattern matching, using POSIX extended regular expression syntax. The first ⟨general text⟩ is a pattern and the second is a string. The command expands to -1 if the pattern is invalid, to 0 if no match is found, and to 1 if a match is found. With the icase option, the matching is case-insensitive. The subcount option sets the size of the table storing the found (sub)patterns; its default is 10. The primitive was introduced in pdfTeX 1.30.0.

4.15.7 \pdflastmatch

\pdflastmatch ⟨integer⟩ (expandable)

The matches found with \pdfmatch are stored in a table. This command returns the entry for match ⟨integer⟩, in the format ⟨position⟩->⟨string⟩; ⟨position⟩ is the position of the match (starting at zero) or -1 if no match was found, and ⟨string⟩ is the matched substring.

Entry 0 contains the match as a whole; the subsequent entries contain submatches corresponding to the subpatterns, up to subcount-1.

If ⟨integer⟩ is less than zero, an error is given.

For instance:

```
\pdfmatch subcount 3 {ab(cd)*ef(gh)(ij)}{abefghij}
\pdflastmatch0 % "0->abefghij"
\pdflastmatch1 % "-1->"
\pdflastmatch2 % "4->gh"
\pdflastmatch3 % "-1->"
```

Entry 1 is empty because no match was found for `cd`, and entry 3 is empty because it exceeds the table's size as set by `subcount`. The primitive was introduced in pdfTEX 1.30.0.

4.16   Numbers

4.16.1   \ifpdfabsnum, \ifpdfabsdim

\ifpdfabsnum   (expandable)
\ifpdfabsdim   (expandable)

This conditional works like the standard `\ifnum` (resp. `\ifdim`), except that it compares absolute values of numbers (dimensions). Although it seems to be a trivial shortcut for a couple of regular `\if` tests, as primitives they are considerably simpler and faster to use. The primitive was introduced in pdfTEX 1.40.0.

4.16.2   \pdfnormaldeviate

\pdfnormaldeviate   (expandable)

Generate a normally-distributed random integer with a mean of 0 and standard deviation 65536. That is, about 68% of the time, the result will be between -65536 and 65536 (one standard deviation away from the mean). About 95% of results will be within two standard deviations, and 99.7% within three. This primitive expands to a list of tokens. The primitive was introduced in pdfTEX 1.30.0.

4.16.3   \pdfuniformdeviate

\pdfuniformdeviate ⟨number⟩   (expandable)

Generate a uniformly-distributed random integer value between 0 (inclusive) and ⟨number⟩ (exclusive). This primitive expands to a list of tokens. The primitive was introduced in pdfTEX 1.30.0.

4.16.4   \pdfrandomseed

\pdfrandomseed   (read-only integer)

You can use `\pdfrandomseed` to query the current seed value, so you can e.g., write the value to the log file. The initial value of the seed is derived from the system time, and is not more than 1 000 999 999 (this ensures that the value can be used with commands like `\count`). The primitive was introduced in pdfTEX 1.30.0.

4.16.5   \pdfsetrandomseed

\pdfsetrandomseed ⟨number⟩

Set the random seed (`\pdfrandomseed`) to a specific value, allowing you to replay sequences of semi-randoms at a later moment. The primitive was introduced in pdfTEX 1.30.0.

4.17   Timekeeping

4.17.1   \pdfelapsedtime

\pdfelapsedtime   (read-only integer)

Return a number that represents the time elapsed from the moment of the start of the run. The elapsed time is returned in 'scaled seconds', meaning seconds divided by 65536, e.g., pdfTEX has run for 50827 'scaled seconds' when this paragraph was typeset. The primitive will never return

a value greater than the highest number available in TeX: if the time exceeds 32767 seconds, the constant value $2^{31} - 1$ will be returned. The primitive was introduced in pdfTeX 1.30.0.

4.17.2 \pdfresettimer

\pdfresettimer

Reset the internal timer so that \pdfelapsedtime starts returning the scaled seconds from 0 again. The primitive was introduced in pdfTeX 1.30.0.

4.18 Files

4.18.1 \pdffiledump

\pdffiledump [offset ⟨integer⟩] [length ⟨integer⟩] ⟨general text⟩   (expandable)

Expands to the dump of the first length ⟨integer⟩ bytes of the file specified by ⟨general text⟩, in uppercase hexadecimal format (same as \pdfescapehex), starting at offset offset ⟨number⟩, or the beginning of the file if offset is zero or not specified. If length is not given, the default is zero, so expands to nothing. Both ⟨integer⟩s must be ≥ 0. For example, the first ten bytes of the source of this manual are 2520706466546558206D. The primitive was introduced in pdfTeX 1.30.0.

4.18.2 \pdffilesize

\pdffilesize ⟨general text⟩   (expandable)

Expands to the size of the file specified by ⟨general text⟩, e.g., 186166 for the source of this manual. The primitive was introduced in pdfTeX 1.30.0.

4.18.3 \pdfmdfivesum file

\pdfmdfivesum [file] ⟨general text⟩   (expandable)

If the keyword file is given, expands to the MD5 checksum of the ⟨general text⟩ in uppercase hexadecimal format (same as \pdfescapehex). Without file, expands to the MD5 of the ⟨general text⟩ taken as a string. For example, the MD5 of the source for this manual is CEE3A702C5A614A5088CD097370FF943. The primitive was introduced in pdfTeX 1.30.0.

4.18.4 \pdffilemoddate

\pdffilemoddate ⟨general text⟩   (expandable)

Expands to the modification date of file specified by ⟨general text⟩, in the same format as for \pdfcreationdate, e.g., D:20240107094514-08'00' for the source of this manual. As of pdfTeX 1.40.20, if the SOURCE_DATE_EPOCH and FORCE_SOURCE_DATE environment variables are both set, \pdffilemoddate returns a value in UTC, ending in Z. The primitive was introduced in pdfTeX 1.30.0.

4.18.5 \input

\input ⟨general text⟩   (expandable)

As of TeX Live 2020, the \input primitive in all TeX engines, including pdfTeX, now also accepts a group-delimited filename argument, as a system-dependent extension, as in \input{foo.tex}. The standard usage of \input with a space/token-delimited filename is completely unchanged.

This group-delimited argument was previously implemented in LuaTeX; now it is available in all engines. ASCII double quote characters (") are removed from the filename, but it is otherwise left unchanged after tokenization.

This extension is unlike most others in that it affects a primitive in standard TEX (\input), rather than being related to a new primitive, command line option, etc. This is allowed because additional methods of recognizing filenames are explicitly mentioned in `tex.web` as acceptable system-dependent extensions.

Incidentally, this does not directly affect LATEX's \input command, as that is a macro redefinition of the standard \input primitive.

The primitive was introduced in pdfTEX 1.40.21.

4.19 Color stack

pdfTEX 1.40.0 introduced color stack support; more generally, stacks of any graphic state.

4.19.1 \pdfcolorstackinit

\pdfcolorstackinit [`page`] [`direct`] ⟨general text⟩   (expandable)

The primitive initializes a new graphic stack and returns its number. The optional `page` keyword instructs pdfTEX to restore the graphic at the beginning of every new page. The optional `direct` keyword has the same effect as for \pdfliteral primitive. The primitive was introduced in pdfTEX 1.40.0.

4.19.2 \pdfcolorstack

\pdfcolorstack ⟨stack number⟩ ⟨stack action⟩ ⟨general text⟩

⟨stack action⟩ → `set` | `push` | `pop` | `current`

The command operates on the stack of a given ⟨stack number⟩. If ⟨stack action⟩ is:

push: the new value in ⟨general text⟩ is inserted at the top of the graphic stack and becomes the current stack value.

pop: the top value is removed from the stack and the new top value becomes the top of the stack.

set: the current value is replaced with ⟨general text⟩, without changing the stack size.

current: the current stack value is return without modifying the stack.

The primitive was introduced in pdfTEX 1.40.0.

4.20 Transformations

Since the content of \pdfliteral is not interpreted, any transformation inserted directly into the content stream, as well as saving and restoring the current transformation matrix, will be unnoticed by pdfTEX's positioning mechanism. As a consequence, links and other annotations (that are formed in PDF on a different layer then the page content) are not affected by such user-defined transformations. pdfTEX 1.40.0 solves this problem with three new primitives.

4.20.1 \pdfsetmatrix

`\pdfsetmatrix`

Affine transformations are normally expressed with six numbers. The first four (no unit) values defining scaling, rotating and skewing, plus two extra dimensions for translation (shifting). Since translation is handled by TeX itself, the `\pdfsetmatrix` primitive expects as an argument a string containing just the first four numbers of the transformation separated by a space and assumes the two position coordinates to be 0.

For example, `\pdfsetmatrix{0.87 -0.5 0.5 0.87}` rotates the current space by 30 degrees, inserting `0.87 -0.5 0.5 0.87 0 0 cm` into the output. The primitive was introduced in pdfTeX 1.40.0.

4.20.2 \pdfsave

`\pdfsave`

The command saves the current transformation by inserting the `q` operator into the content stream. The primitive was introduced in pdfTeX 1.40.0.

4.20.3 \pdfrestore

`\pdfrestore`

The command restores previously saved transformation by inserting the `Q` operator into the content stream. `\pdfsave` and `\pdfrestore` pairs should always be properly nested and should start and end at the same group level. The primitive was introduced in pdfTeX 1.40.0.

4.21 Macro programming

Expansion-related commands.

4.21.1 \expanded

`\expanded` ⟨tokens⟩   (expandable)

Expands ⟨tokens⟩ in exactly the same way as `\message`. In contrast to `\edef`, macro parameter characters do not need to be doubled. `\protected` macros (q.v.) are not expanded. The primitive was introduced in pdfTeX 1.40.20.

4.21.2 \ifincsname

`\ifincsname`   (expandable)

This conditional is true if evaluated inside `\csname ... \endcsname`, and false otherwise.

4.21.3 \ifpdfprimitive

`\ifpdfprimitive` ⟨control sequence⟩   (expandable)

This conditional checks if the following control sequence has its primitive meaning. If so, `\ifpdfprimitive` is true. In any other case (redefined, made `\undefined`, has never been primitive), it is false. The primitive was introduced in pdfTeX 1.40.0.

4.21.4 \partokencontext

\partokencontext ⟨number⟩

The *par-token* (i.e., the token with the name given by \partokenname, which is \par by default; see \partokenname, next) is inserted into the input stream in different places, according to the \partokencontext value. This value can be:

0 (the default): *par-token* is inserted at empty lines (more exactly, when a token category 5 is seen in state N, reading a line); before \end, \vskip, \hrule, \unvbox, and \halign, if TeX is in horizontal mode when those commands are seen; and in various error recovery situations. These are the standard cases in TeX.

1: *par-token* is inserted as above, and also at the end of \vbox, \vtop, and \vcenter, if TeX is in horizontal mode at the time.

2: *par-token* is inserted as above, and also at the end of \insert, \vadjust, \output, \noalign, and items of \valign, again if TeX is in horizontal mode at the time.

With the default \partokencontext=0, TeX behaves in its normal way: the situations in cases 1 and 2 are processed by a direct call of *end-paragraph* routine, with no emitted *par-token*.

If \partokencontext=1 then TeX inserts the *par-token* in additional cases: when vertical boxes are completed but horizontal mode is not finished. Since vboxes are not uncommonly inserted directly by users, with horizontal mode material, this allows macro programmers to control all such boxes being finished by a *par-token*. An example:

```
\partokenname\_mypar
\partokencontext=1
\def\_mypar{Hi there!\endgraf}
\vbox{Vbox text.}
```

This will output "Hi there!" after "Vbox text.".

Finally, with \partokencontext=2, all cases where classical TeX uses the direct *end-paragraph* routine are changed to emit the *par-token* instead. In contrast to case 1, these commands are rarely invoked directly by users with horizontal mode material.

The setting of the register \partokencontext is local.

The primitive was introduced in pdfTeX 1.40.24.

4.21.5 \partokenname

\partokenname ⟨control sequence⟩

TeX internally inserts a control sequence, named \par by default, into the input stream at empty lines, the end of vboxes, and various other places (see \partokencontext, above). Let's call this control sequence the *par-token*.

Executing \partokenname⟨control sequence⟩ changes the name of the *par-token* from \par to the given ⟨control sequence⟩. The setting performed by \partokenname is global.

This makes it possible to release the name \par to the "user's name space", i.e., after using \partokenname, users can define and use \par as they wish without changing the behavior of anything internal to TeX. For example:

```
\catcode`\_=11
\partokenname\_mypar % use \_mypar at user level
\let\_mypar=\par      % make \_mypar equivalent to built-in \par
%
\def\par{some random text} % redefine \par
%
```

```
Hello world.

Goodbye.
\end
```

This will not output "`some random text`" (the definition of `\par`), due to the `\partokenname` setting.

By default, the meaning of the *par-token* is to end a paragraph (also named as `\endgraf` in the plain TEX format). It can be changed as usual with, for example, `\def`. Naturally, it is the control sequence name given to `\partokenname` that must be redefined. Continuing the previous example (prior to the `\end`):

```
\def\_mypar{Hi there!\endgraf}
Paragraph one.

Paragraph two.\let\_mypar=\endgraf
```

This will output "Hi there!" after "Paragraph one.", before ending the paragraph.

Another behavior of the *par-token* built into TEX is that macros not defined as `\long` cause the error "runaway argument" if the *par-token* is scanned as a parameter. After `\setpartokenname`, it will be the new control sequence name that triggers this error, not `\par`. For instance (still continuing the same example):

```
\def\amac#1{}
\amac{long test, no error: \par}
\amac{long test, gives error: \_mypar}
```

The primitive was introduced in pdfTEX 1.40.24.

4.21.6 `\pdfprimitive`

`\pdfprimitive` ⟨control sequence⟩

This command executes the primitive meaning of the following control sequence, regardless of whether the control sequence has been redefined or made undefined. If the primitive was expandable, `\pdfprimitive` expands also. On the other hand, if the following control sequence never was a primitive, nothing happens and no error is raised. (In some versions of pdfTEX prior to 1.40.19, an error was wrongly given.) The primitive was introduced in pdfTEX 1.40.0.

4.22 Typesetting

4.22.1 `\pdfinsertht`

`\pdfinsertht` ⟨integer⟩   (expandable)

If ⟨integer⟩ is the number of an insertion class, this command returns the current height of the corresponding box. For instance, the following returns `12pt` in plain TEX:

```
Abc\footnote*{Whatever.}\par
\pdfinsertht\footins
```

4.22.2 \pdflastxpos, \pdflastypos

\pdflastxpos   (read-only integer)
\pdflastypos   (read-only integer)

This primitive returns an integer number representing the absolute x resp. y coordinate of the last point marked by \pdfsavepos. The unit is scaled points (sp).

4.22.3 \pdfsavepos

\pdfsavepos (h, v, m)

This primitive marks the current absolute (x, y) position on the media, with the reference point in the lower left corner. It is active only during page shipout, when the final page is assembled. The position coordinates can then be retrieved by the \pdflastxpos and \pdflastypos primitives, and e.g., written out to some auxiliary file. The coordinates can be used only after the current \shipout has been finalized, therefore normally two pdfTeX runs are required to utilize these primitives. Starting with pdfTeX 1.40.0, this mechanism can also be used while running in DVI mode.

4.22.4 \quitvmode

\quitvmode

The primitive instructs pdfTeX to quit vertical mode and start typesetting a paragraph. Thus, \quitvmode has the same basic effect as the \leavevmode macro from plain.tex. However, \leavevmode expands the \everypar token list, which may or may not be desired. \quitvmode does not expand \everypar. The primitive was introduced in pdfTeX 1.21a.

4.22.5 \vadjust

\vadjust [⟨pre spec⟩] ⟨filler⟩ {⟨vertical mode material⟩ } (h, m)

In pdfTeX, the \vadjust primitive supports an additional optional qualifier ⟨pre spec⟩, which is simply the string pre, to the original TeX. If no pre is given, \vadjust behaves exactly as the original (see *The TeXbook*, p. 281): it appends an adjustment item created from ⟨vertical mode material⟩ to the current list *after* the line in which \vadjust appears. In contrast, with the qualifier pre, the adjustment item is put *before* the line in which \vadjust pre appears.

4.23   Tracing

4.23.1   \showstream

\showstream   (integer)

If this primitive parameter has a value corresponding to an open output stream (which has been opened with \openout), then any \show, \showthe, \showbox or \showlists commands do not write output to the terminal, but instead write only to the referenced output stream, as if they were written with \immediate\write.

For example:

```
\newwrite\myoutstream
\immediate\openout\myoutstream="infofile"
\showstream=\myoutstream
% From now on, \show... commands are redirected to "infofile.tex".
\show\TeX
%
```

```
\showstream=-1
% -1 is never a open file and therefore restores
% normal \show... behavior.
\immediate\closeout\myoutstream
```

This example would not generate any special output to the terminal or log file (except for any logging done by \newwrite. It writes this text to `infofile.tex`, including the initial blank line, since that is what \show does:

```
> \TeX=macro:
->T\kern -.1667em\lower .5ex\hbox {E}\kern -.125emX
```

This primitive is available, with identical behavior, in all TEX engines except the original TEX and ε-TEX, where \showstream remains undefined. The primitive was introduced in pdfTEX 1.40.24.

4.23.2 \tracinglostchars

\tracinglostchars (integer)

This primitive parameter has always been part of TEX, and its operation with values ≤ 2 is unchanged. In addition, if its value is ≥ 3, then "Missing character" reports become full errors (ordinarily they are only logged), and the character code is reported in hex. For example:

```
\tracinglostchars=3
\font\x=logo10 \x \char99 \end
```

will result in this error message:

```
! Missing character: There is no c ("63) in font logo10.
```

(The `logo10` font only defines the capital letters used in the METAFONT and MetaPost logos, so there is no lowercase.)

This new behavior is essentially the same in all TEX engines except the original TEX and ε-TEX, where the behavior of \tracinglostchars remains unchanged.

The primitive was introduced in pdfTEX 1.40.22.

4.23.3 \tracingstacklevels

\tracingstacklevels (integer)

If this primitive parameter is > 0, and \tracingmacros > 0, macro expansion logging is truncated at the specified depth. Also, and more importantly, each relevant log line is given a prefix beginning with ~, either followed by a . character for each expansion level or another ~ if the expansion was truncated. For example:

```
\tracingmacros=1      % so macro expansion is logged at all
\tracingstacklevels=2 % cut off at level 2
\def\a#1{\relax}      % argument to show parameter logging is affected too
\def\b#1{\a{#1}}
\b1
```

logs the following:

```
~.\b #1->\a {#1}
#1<-1
~~\a
```

Thus, the expansion of \b is logged normally, with the addition of the ~. prefix. The expansion of \a is truncated (level 2), hence neither the parameters nor body expansion are shown.

Furthermore, an \input file counts as an expansion level, and the input filename is logged. So, if we add this to our example above:

```
\input anotherfile
```

where anotherfile.tex simply contains \b2, the log will get:

```
~.INPUT anotherfile.tex
~~\b
~~\a
```

Now the \b expansion is not logged either, since the expansion depth is higher than the \tracingstacklevels value.

The intended use of \tracingstacklevels is not so much to truncate logging as to indicate expansion levels for detailed debugging. Thus normally it would be set to a large number (\maxdimen), so that everything is fully logged, with the addition of the expansion level indication with the number of dots in the prefix.

This primitive is available, with identical behavior, in all TeX engines except the original TeX and ε-TeX, where \tracingstacklevels remains undefined. The primitive was introduced in pdfTeX 1.40.22.

4.24   pdfTeX execution environment

4.24.1   \pdfdraftmode

\pdfdraftmode   (integer)

When set to 1 (or set by the command-line switch -draftmode), pdfTeX doesn't write the output PDF file and doesn't read any images, but does do everything else (including writing auxiliary files), thus speeding up compilations when you know you need an extra run but don't care about the output, e.g., just to get cross-references started. If specified, the parameter must appear before any data is written to the PDF output. The primitive was introduced in pdfTeX 1.40.0.

4.24.2   \pdfshellescape

\pdfshellescape   (read-only integer)

This primitive is 1 if \write18 is enabled, 2 if its operation is restricted to known-safe programs, and 0 otherwise. The primitive was introduced in pdfTeX 1.30.0.

4.24.3   \pdftexbanner

\pdftexbanner   (expandable)

Returns the pdfTeX banner message, e.g., for the version used here: `This is pdfTeX, Version 3.141592653-2.6-1.40.25 (TeX Live 2023) kpathsea version 6.3.5`. The primitive was introduced in pdfTeX 1.20a.

4.24.4 \pdftexrevision

`\pdftexrevision`   (expandable)

Returns the revision number of pdfTeX, e.g., for pdfTeX version 1.40.25 (used to produce this document), it returns the number 25.

4.24.5 \pdftexversion

`\pdftexversion`   (read-only integer)

Returns the version of pdfTeX multiplied by 100, excluding the tertiary version number, e.g., for pdfTeX version 1.40.25 (used to produce this document), it returns 140.

Chapter 5

Graphics

pdfTeX supports inclusion of pictures in PNG, JPEG, JBIG2, and PDF format; a few differences between these are discussed below. The most common technique with TeX, the inclusion of EPS figures, is not (and cannot be) supported; PDF inclusion is usually the easiest replacement. EPS files can be converted to PDF by Ghostscript, Adobe Distiller or other PostScript-to-PDF converters.

pdfTeX allows inserting selected pages from PDF files, carry their own fonts, graphics, and pixel images, into a document. The cover page of this manual is an example of such an insert, being a one-page document generated by pdfTeX.

By default pdfTeX takes the BoundingBox of a PDF file from its CropBox if available, otherwise from its MediaBox. This can be influenced by the ⟨pdf box spec⟩ option to the `\pdfximage` primitive, or by setting the `\pdfpagebox` or `\pdfforcepagebox` primitives to a value corresponding to the desired box type.

To get the right BoundingBox from a EPS file, before converting to PDF, it is necessary to transform the EPS file so that the start point is at the (0,0) coordinate and the page size is set exactly corresponding to the BoundingBox. A Perl script (epstopdf) for this purpose and the PStoPDF program that comes with Ghostscript can do a similar job. (Concerning this conversion, they can handle complete directories, remove some garbage from files, takes precautions against duplicate conversion, etc.)

The lossless compressing PNG format is useful for embedding crisp pixel graphics (e.g., line scans, screenshots). As of pdfTeX 1.30.0, the alpha channel of PNG images is processed if available; this allows embedding of images with simple transparency. The PNG format does not support the CMYK color model, which is sometimes required for print media (this often can be replaced by four component JPEG in high quality or lossless compression mode). Photos in PNG format have rather poor compression; here the JPEG format is preferable.

Embedding PNG images in the general case requires pdfTeX to uncompress the pixel array and to re-compress it to the PDF requirements; this may take a noticeable amount of time. Since pdfTeX 1.30.0 there is now a fast PNG embedding mode that goes without uncompressing; the image data are directly copied into the PDF stream, resulting in a much higher embedding speed. This direct mode is only activated if the image array structure of the PNG file is compatible with the PDF image structure; e.g., an interlaced PNG image requires uncompressing to re-arrange the image lines. Luckily, it seems that in practice PNG files commonly allow fast copying. The use of gamma correction disables fast copying, as it requires calculations with individual pixels. Whether the fast copy mode is used for a PNG image can be seen from the log file, which then shows the string '(PNG copy)' after the PNG file name.

The JPEG format is normally used in lossy mode; then it's ideal for embedding photos. It's not recommended for crisp images from synthetic sources with a limited amount of colors. Both JFIF and EXIF are supported for additional information.

The JBIG2 format works only for bi-tonal (black and white) pixel images like scanned line and text documents, but for these it has typically a much higher compression ratio than the other two pixel image formats. The JBIG2 format is part of the PDF standard since version 1.5; native JBIG2 image inclusion is available in pdfTeX since version 1.40.0. A JBIG2 file might contain many images, which gives an even better compression ratio than with a single image per file, as JBIG2 encoders can exploit similarities between bit patterns over several images. Encoders for JBIG2 can operate in lossy as well as lossless modes. A free JBIG2 encoder is available: see `https://github.com/agl/jbig2enc` (perhaps there are others).

Other options for graphics in pdfTeX are:

LaTeX picture mode Since this is implemented simply in terms of font characters, it works in exactly the same way as usual.

Xy-pic If the PostScript backend is not requested, Xy-pic uses its own Type 1 fonts, and needs no special attention.

tpic The tpic `\special` commands (used in some macro packages) can be redefined to produce literal PDF, using some macros written by Hans Hagen.

MetaPost Although the output of MetaPost is PostScript, it is in a highly simplified form, and a MetaPost to PDF conversion (mptopdf, available as `supp-pdf.tex` and `supp-pdf.mkii`, written by Hans Hagen and Tanmoy Bhattacharya) is implemented as a set of macros which reads MetaPost output and supports all of its features.

The inclusion of raw PostScript commands, a technique utilized by for instance the `pstricks` package, cannot directly be supported. Although PDF is a direct descendant of PostScript, it lacks any programming language commands, and cannot deal with arbitrary PostScript.

Chapter 6

Additional PDF keys

This section is based on the manual on keys written by Martin Schröder.

A PDF document should contain only the structures and attributes defined in the PDF specification. However, the specification allows applications to insert additional keys, provided they follow certain rules.

The most important rule is that developers have to register with Adobe prefixes for the keys they want to insert. Hans Hagen has registered the prefix `PTEX` for pdfTEX.

pdfTEX generates an XObject for every included PDF image. The dictionary of this object contains these additional keys:

key	type	meaning
PTEX.FileName	string	The name of the included file as seen by pdfTEX.
PTEX.InfoDict	dictionary	The `InfoDict` of the included PDF (an indirect object).
PTEX.PageNumber	integer	The page number of the included file.

The PDF reference manual says: "Although viewer applications can store custom metadata in the document information dictionary, it is inappropriate to store private content or structural information there; such information should be stored in the document catalog instead."

Although it would seem more natural to put this information in the document information dictionary, we have to obey the rules laid down in the PDF reference manual. The following key ends up in the document catalog.

key	type	meaning
PTEX.Fullbanner	string	The full version of the pdfTEX binary that produced the file as displayed by `pdftex --version`, a.k.a. `\pdftexbanner` (`This is pdfTeX, Version 3.141592653-2.6-1.40.25 (TeX Live 2023) kpathsea version 6.3.5`). This is necessary because the string in the `Producer` key in the info dictionary is rather short, namely `pdfTeX-1.40.25`.

Any or all of these keys can be suppressed with the `\pdfsuppressptexinfo` primitive, described in section 4.2.

Chapter 7

Character translation

Characters that are input to pdfTEX are subject to optional TEX character translation (TCX) under control of a TCX file. The TCX maps the input character codes (e.g., from `\input` or `\read`) to the character codes as seen by pdfTEX. This mapping takes place before the characters enter pdfTEX's 'mouth'. If no TCX file is read, the input characters enter pdfTEX directly; no mapping is done.

TCX files consist of lines each containing one or two integer numbers in the range 0..255, either in decimal or hex notation. A comment sign `%` in a TCX line starts a comment until the end of line. The first number in each line is for matching the input character code, the second, optional number is the corresponding TEX character code. If a line contains only one number, characters with this code enter pdfTEX unchanged; no mapping is done.

TCX mapping also influences pdfTEX output streams for `\message` and `\write`. Without TCX mapping, only characters that are within the range 32..126 are flagged as 'printable', meaning that these characters are output directly by `\message` and `\write` primitives. Characters outside the range 32..126 are instead output in escaped form, e.g., as `^^A` for a character with code `0x01`. When a character code is mentioned in the 2nd column of the TCX file, or as the only value in a line, it is flagged as 'printable'. During `\message` and `\write`, output characters are mapped in reverse direction: they are looked up in the 2nd column of the TCX file and the corresponding values from the 1st column are output. Again, if a pdfTEX character code is found as the only number in a line, no mapping is done. Mentioning a character code as the only number on a line has the sole purpose to flag this code 'printable'.

The characters output into the PDF file, e.g., by the `\pdfliteral` or `\special` primitives, are not subject to TCX output remapping.

Beware: Character translation interferes with the encTEX primitives; to avoid surprises, don't use encTEX and TCX mapping at the same time. Further details about TCX file loading can be found in the Web2C manual, `https://tug.org/texinfohtml/web2c.html#TCX-files`.

Nowadays TCX files are rarely used. The `-8bit` command line option makes all characters printable, when that is desired.

Appendix A

Installation

This section describes the steps needed to get pdfTeX running on a system where pdfTeX is not yet installed. Nowadays all known TeX distributions, such as TeX Live, MiKTeX, and MacTeX, include pdfTeX already. For example, the TeX Live distribution comes with pdfTeX versions for many Unix, Windows, and Mac systems; more information can be found at `https://tug.org/texlive`. When you use any of these distributions, you don't need to bother with the pdfTeX installation procedure in this chapter.

If there is no precompiled pdfTeX binary for your system, or the version coming with a distribution is not the current one and you would like to try out a fresh pdfTeX immediately, you will need to build pdfTeX from sources; read on. You should already have a working TeX system, e.g., TeX Live, into which the freshly compiled pdfTeX will be integrated. (Numerous support files are required in order to have a working binary.) Note that the installation description in this manual is Web2C-specific.

A.1 Getting sources and binaries

The primary home page is `https://www.pdftex.org`, where you also find bug tracking information. Development sources are at `svn://tug.org/pdftex/branches/stable`. Precompiled binaries for some platforms are available in subdirectories below `https://ctan.org/tex-archive/systems`, although these are unlikely to be up to date. Your best bet is to use a current TeX distribution.

A.2 Compiling

Given a directory where the sources have been downloaded or checked out, some working on a Unix system the following steps are needed to compile pdfTeX:

```
cd source  # from branches/stable!
./build-pdftex.sh
```

The binary `pdftex` will hopefully then end up in the subdirectory `build-pdftex/texk/web2c`. The additional utilities `pdftosrc` and `ttf2afm` are in the pdfTeX source repository, but are not built by default. If they are needed, edit `build-pdftex.sh` to `--enable` them.

For pdfTeX maintainers: a sibling script to `build-pdftex.sh` is in the repository, namely `sync-pdftex.sh`, which syncs changes from a TeX Live source repository to a pdfTeX source repository. Read and understand the script before using it!

A.3   Placing files

The next step is to put the newly-compiled `pdftex` (and, if enabled, `pdftosrc` and/or `ttf2afm`) binaries into the installed distribution. E.g., for a typical TeX Live system: `/usr/local/texlive/2023/bin/x86_64-linux`, replacing `x86_64-linux` with the appropriate platform name. (As always, it is sensible to rename the installed `pdftex` binary first and not just overwrite it.)

If you're doing this into a live hierarchy, run `fmtutil-sys refresh` afterwards, so that all formats are regenerated system-wide with the new `pdftex` binary.

If what you want is to test a new version of `pdftex`, and not replace the installed version, the best approach is to copy it into the binary directory as, say, `pdftex.new`. The different executable name can sometimes change the files found and other behavior, but usually this won't matter.

pdfTeX uses the Kpathsea library to search for supporting files; many variables and configuration files can come into play. See the Kpathsea manual: `https://tug.org/kpathsea`.

A.4   Configuration of pdfTeX

As opposed to TeX with its DVI output, the pdfTeX program does not have a separate postprocessing stage to transform the TeX input into final PDF. As a consequence, all data needed for building the PDF must be available during the pdfTeX run, in particular information on media dimensions and offsets, graphics files for embedding, and font information (font files, encodings).

When TeX builds a page, it places items relative to the (1in,1in) offset from the top left page corner (the DVI reference point). Separate DVI postprocessors allow specifying the paper size (e.g., A4 or letter), so that this reference point is moved to the correct position on the paper, and the text ends up at the right place.

In PDF, the paper dimensions are part of the page definition, and pdfTeX therefore requires that they be defined at the beginning of the pdfTeX run. As with pages described by PostScript, the PDF reference point is in the lower left corner.

Formerly, these dimensions and other pdfTeX parameters were read in from a configuration file named `pdftex.cfg`, which had a special (non-TeX) format, at the start of processing. Nowadays such a file is ignored by pdfTeX. Instead, the page dimensions and offsets, as well as many other parameters, can be set by pdfTeX primitives during the pdfTeX format building process, so that the settings are dumped into the fresh format and consequently will be used when pdfTeX is later called with that format. All settings from the format can still be overridden during a pdfTeX run by using the same primitives. This configuration concept is a more unified approach, as it avoids a configuration file with a special format.

A list of pdfTeX primitives likely relevant to setting up the pdfTeX engine is given in the table below. All primitives are described in detail in previous sections.

```
% tex-ini-files 2016-04-15: pdftexconfig.tex

% Load shared (PDF) settings in pdfTeX

% Enable PDF output
\pdfoutput            = 1

% Paper size: dimensions given in absolute terms
\pdfpageheight        = 11 true in
\pdfpagewidth         = 8.5 true in

% Enable PDF 1.5 output and thus more compression
\pdfminorversion      = 5
\pdfobjcompresslevel  = 2

% Low-level settings unlikely ever to need to change
\pdfcompresslevel     = 9
\pdfdecimaldigits     = 3
\pdfpkresolution      = 600
\pdfhorigin           = 1 true in
\pdfvorigin           = 1 true in
```

Figure A.1: pdfTeX configuration file for TeX Live (`pdftexconfig.tex`).

primitive name	type	default	comment
\pdfoutput	integer	0	DVI by default
\pdfadjustspacing	integer	0	off
\pdfcompresslevel	integer	9	best
\pdfobjcompresslevel	integer	0	no object streams
\pdfdecimaldigits	integer	3	
\pdfimageresolution	integer	72	dpi
\pdfpkresolution	integer	0	
\pdfpkmode	tokens	empty	mode per `mktex.cnf`
\pdfuniqueresname	integer	0	
\pdfprotrudechars	integer	0	
\pdfgentounicode	integer	0	
\pdfmajorversion	integer	1	output PDF 1.x
\pdfminorversion	integer	4	specifically PDF 1.4
\pdfpagebox	integer	0	
\pdfforcepagebox	integer	0	
\pdfinclusionerrorlevel	integer	0	
\pdfhorigin	dimension	1in	
\pdfvorigin	dimension	1in	
\pdfpagewidth	dimension	0pt	
\pdfpageheight	dimension	0pt	
\pdflinkmargin	dimension	0pt	
\pdfdestmargin	dimension	0pt	
\pdfthreadmargin	dimension	0pt	
\pdfmapfile	text	`pdftex.map`	not dumped

Figure A.1 shows a recent pdfTeX configuration file (`pdftexconfig.tex`) from TeX Live (part of the `tex-ini-files` package), overriding some of these settings. It is read when a format is built. It enables PDF output, sets paper dimensions and the default pixel density for PK font inclusion. The default values are chosen so that pdfTeX often can be used (e.g., in `-ini` mode) without setting any additional parameters.

Independent of whether such a configuration file is read, the first action in a pdfTeX run is

```
% Thomas Esser, 1998. public domain.
\input etex.src
\dump
\endinput
```

Figure A.2: The `etex.ini` file to dump the plain ε-TeX format with DVI output.

```
% Thomas Esser, 1998. public domain.
% This is used for pdftex and pdfetex, which are now identical: both
% with e-TeX extensions, both with pdf output.
\input pdftexconfig.tex
\input etex.src
\input pdftexmagfix.tex
\dump
\endinput
```

Figure A.3: The `pdfetex.ini` file to dump plain ε-TeX with PDF output.

reading the global Web2C configuration file (`texmf.cnf`), which is common to all programs in the Web2C system. This file mainly defines file search paths, the memory layout (e.g., string pool and hash size), and a few other general parameters.

A.5 Creating format files

The pdfTeX engine supports building separate formats for either DVI or PDF output in the same way as the classical TeX engine does for DVI. Format generation (and other `initex` features) is enabled by the `-ini` option. The default mode (DVI or PDF) can be chosen either on the command line by setting the option `-output-format` to `dvi` or `pdf`, or by setting the `\pdfoutput` parameter. The format file inherits this setting, so that a later invocation of pdfTeX with this format starts in the preselected mode (which can still be overridden). A format file can be read in only by the engine that has generated it; a format incompatible with an engine leads to a fatal error.

It is customary to package the configuration and macro file input into a `.ini` file. E.g., the file `etex.ini` in figure A.2 is for generating an ε-TeX format with DVI output. It has been traditional for many years to generate `etex.fmt` with pdfTeX rather than the original ε-TeX, because pdfTeX contains useful additional programming and other non-PDF-related features.

The `pdfetex.ini` file figure A.3 shows the corresponding format with PDF output by default; this is what creates the format file read when `pdftex` is normally invoked.

The corresponding pdfTeX invocations for format generation are:

```
pdftex -ini *etex.ini
pdftex -ini *pdfetex.ini
```

These calls produce format files `etex.fmt`, `pdfetex.fmt`, as the default format file name is taken from the input file name. You can override this with the `-jobname` option. The asterisk `*` before the file name is an unusual flag, only used in `-ini` mode, which causes the pdfTeX engine to enable ε-TeX's features.

To reiterate, the distribution (TeX Live, MiKTeX) usually takes care of format (re)generation. The above is if you need to do such things manually for testing, debugging, development, etc.

Incidentally, as of pdfTeX 1.40.21 (TeX Live 2020), `.fmt` files are compressed with `zlib`. This makes for a considerable savings in space, and consequently in time.

A.6   Testing the installation

When everything is set up, you can test the installation. A simple test of plain pdfTEX is:

```
pdftex story \\end
```

This should typeset the famous one-page short story by A.U. Thor, generating a PDF file.

A more thorough and descriptive test is the plain TEX test file `samplepdf.tex`, available in the distribution in the `samplepdftex/` directory. Process this file by typing:

```
pdftex samplepdf
```

If the installation is ok, this should produce a file called `samplepdf.pdf`. The file `samplepdf.tex` is also a good place to look for examples of how to use pdfTEX's primitives.

A.7   Common problems

The most common problem with installations is that pdfTEX complains that some file cannot be found. In such cases, first make sure that all TEX-related environment variables are unset. For detailed debugging, set the environment variable `KPATHSEA_DEBUG=255` before running pdfTEX or use the option `-kpathsea-debug 255`. More options can be found in the Web2C documentation.

Variables in `texmf.cnf` can be overwritten by environment variables. Here are some of the most common problems you can encounter when getting started:

- ```
  I can't find the format file 'pdftex.fmt'!
  I can't find the format file 'pdflatex.fmt'!
  ```

  The format file is not created (see above how to do that) or is not properly placed. Make sure that `TEXFORMATS` in `texmf.cnf` contains the path to `pdftex.fmt` or `pdflatex.fmt`.

- ```
  Fatal format file error; I'm stymied
  ```

  This typically appears if you forgot to regenerate the `.fmt` files after installing a new version of the pdfTEX binary. The first line tells by which engine the offending format was generated.

- pdfTEX cannot find one or more map files (`*.map`), encoding vectors (`*.enc`), virtual fonts, Type 1 fonts, TrueType or OpenType fonts, or some image file.

  Make sure that the required file exists and the corresponding variable in `texmf.cnf` contains a path to the file.

  When you have installed new fonts, and your PDF viewer complains about missing fonts, you should take a look at the log file produced by pdfTEX. Missing fonts, map files, encoding vectors as well as missing characters (glyphs) are reported there.

For more help resources of all kinds, see `https://tug.org/begin`.

Appendix B

Formal syntax specification

This appendix formally specifies the pdfTeX-specific extensions to the TeX macro programming language. Most primitive names are prefixed by 'pdf'. General definitions and syntax rules follow after the list of primitives.

Two new units of measure were introduced in pdfTeX 1.30.0: the new Didot (1nd = 0.375mm) and the new Cicero (1nc = 12nd). The former was proposed by ISO in 1975.

B.1 Integer registers

\efcode ⟨font⟩ ⟨8-bit number⟩ (integer)
\knaccode ⟨font⟩ ⟨8-bit number⟩ (integer)
\knbccode ⟨font⟩ ⟨8-bit number⟩ (integer)
\knbscode ⟨font⟩ ⟨8-bit number⟩ (integer)
\lpcode ⟨font⟩ ⟨8-bit number⟩ (integer)
\pdfadjustinterwordglue (integer)
\pdfadjustspacing (integer)
\pdfappendkern (integer)
\pdfcompresslevel (integer)
\pdfdecimaldigits (integer)
\pdfdraftmode (integer)
\pdfforcepagebox (integer)
\pdfgamma (integer)
\pdfgentounicode (integer)
\pdfimageapplygamma (integer)
\pdfimagegamma (integer)
\pdfimagehicolor (integer)
\pdfimageresolution (integer)
\pdfinclusioncopyfonts (integer)
\pdfinclusionerrorlevel (integer)
\pdfinfoomitdate (integer)
\pdfmajorversion (integer)
\pdfminorversion (integer)
\pdfmovechars (integer)
\pdfobjcompresslevel (integer)

`\pdfomitcharset` (integer)
`\pdfomitinfodict` (integer)
`\pdfomitprocset` (integer)
`\pdfoutput` (integer)
`\pdfpagebox` (integer)
`\pdfpkresolution` (integer)
`\pdfprependkern` (integer)
`\pdfprotrudechars` (integer)
`\pdfsuppressptexinfo` (integer)
`\pdfsuppresswarningdupdest` (integer)
`\pdfsuppresswarningdupmap` (integer)
`\pdfsuppresswarningpagegroup` (integer)
`\pdftracingfonts` (integer)
`\pdfuniqueresname` (integer)
`\rpcode` ⟨font⟩ ⟨8-bit number⟩ (integer)
`\shbscode` ⟨font⟩ ⟨8-bit number⟩ (integer)
`\showstream` (integer)
`\stbscode` ⟨font⟩ ⟨8-bit number⟩ (integer)
`\tagcode` ⟨font⟩ ⟨8-bit number⟩ (integer)
`\tracinglostchars` (integer)
`\tracingstacklevels` (integer)

B.2   Read-only integers

`\pdfelapsedtime` (read-only integer)
`\pdflastannot` (read-only integer)
`\pdflastlink` (read-only integer)
`\pdflastobj` (read-only integer)
`\pdflastxform` (read-only integer)
`\pdflastximage` (read-only integer)
`\pdflastximagecolordepth` (read-only integer)
`\pdflastximagepages` (read-only integer)
`\pdflastxpos` (read-only integer)
`\pdflastypos` (read-only integer)
`\pdfrandomseed` (read-only integer)
`\pdfretval` (read-only integer)
`\pdfshellescape` (read-only integer)
`\pdftexversion` (read-only integer)

B.3　Dimen registers

`\pdfdestmargin`　(dimen)
`\pdfeachlinedepth`　(dimen)
`\pdfeachlineheight`　(dimen)
`\pdffirstlineheight`　(dimen)
`\pdfhorigin`　(dimen)
`\pdfignoreddimen`　(dimen)
`\pdflastlinedepth`　(dimen)
`\pdflinkmargin`　(dimen)
`\pdfpageheight`　(dimen)
`\pdfpagewidth`　(dimen)
`\pdfpxdimen`　(dimen)
`\pdfthreadmargin`　(dimen)
`\pdfvorigin`　(dimen)

B.4　Token registers

`\pdfpageattr`　(tokens)
`\pdfpageresources`　(tokens)
`\pdfpagesattr`　(tokens)
`\pdfpkmode`　(tokens)

B.5　Expandable commands

`\expanded` ⟨tokens⟩　(expandable)
`\ifincsname`　(expandable)
`\ifpdfabsdim`　(expandable)
`\ifpdfabsnum`　(expandable)
`\ifpdfprimitive` ⟨control sequence⟩　(expandable)
`\input` ⟨general text⟩　(expandable)
`\leftmarginkern` ⟨box number⟩　(expandable)
`\pdfcolorstackinit` [`page`] [`direct`] ⟨general text⟩　(expandable)
`\pdfcreationdate`　(expandable)
`\pdfescapehex` ⟨general text⟩　(expandable)
`\pdfescapename` ⟨general text⟩　(expandable)
`\pdfescapestring` ⟨general text⟩　(expandable)
`\pdffiledump` [`offset` ⟨integer⟩] [`length` ⟨integer⟩] ⟨general text⟩　(expandable)
`\pdffilemoddate` ⟨general text⟩　(expandable)
`\pdffilesize` ⟨general text⟩　(expandable)
`\pdffontname` ⟨font⟩　(expandable)
`\pdffontobjnum` ⟨font⟩　(expandable)
`\pdffontsize` ⟨font⟩　(expandable)
`\pdfincludechars` ⟨font⟩ ⟨general text⟩　(expandable)
`\pdfinsertht` ⟨integer⟩　(expandable)

`\pdflastmatch` ⟨integer⟩ (expandable)
`\pdfmatch` [`icase`] [`subcount` ⟨integer⟩] ⟨general text⟩ ⟨general text⟩ (expandable)
`\pdfmdfivesum` [`file`] ⟨general text⟩ (expandable)
`\pdfnormaldeviate` (expandable)
`\pdfpageref` ⟨page number⟩ (expandable)
`\pdfstrcmp` ⟨general text⟩ ⟨general text⟩ (expandable)
`\pdftexbanner` (expandable)
`\pdftexrevision` (expandable)
`\pdfunescapehex` ⟨general text⟩ (expandable)
`\pdfuniformdeviate` ⟨number⟩ (expandable)
`\pdfxformname` ⟨object number⟩ (expandable)
`\pdfximagebbox` ⟨integer⟩ ⟨integer⟩ (expandable)
`\rightmarginkern` ⟨box number⟩ (expandable)

B.6   General commands

`\letterspacefont` ⟨control sequence⟩ ⟨font⟩ ⟨integer⟩
`\partokencontext` ⟨number⟩
`\partokenname` ⟨control sequence⟩
`\pdfannot` ⟨annot type spec⟩ (h, v, m)
`\pdfcatalog` ⟨general text⟩ [`openaction` ⟨action spec⟩]
`\pdfcolorstack` ⟨stack number⟩ ⟨stack action⟩ ⟨general text⟩
`\pdfcopyfont` ⟨control sequence⟩ ⟨font⟩
`\pdfdest` ⟨dest spec⟩ (h, v, m)
`\pdfendlink` (h, m)
`\pdfendthread` (v, m)
`\pdffakespace`
`\pdffontattr` ⟨font⟩ ⟨general text⟩
`\pdffontexpand` ⟨font⟩ ⟨stretch⟩ ⟨shrink⟩ ⟨step⟩ [`autoexpand`]
`\pdfglyphtounicode` ⟨general text⟩ ⟨general text⟩
`\pdfinfo` ⟨general text⟩
`\pdfinterwordspaceoff`
`\pdfinterwordspaceon`
`\pdfliteral` [`shipout`] [`direct` | `page`] ⟨general text⟩ (h, v, m)
`\pdfmapfile` ⟨map filename⟩
`\pdfmapline` ⟨map spec⟩
`\pdfnames` ⟨general text⟩
`\pdfnobuiltintounicode` ⟨font⟩
`\pdfnoligatures` ⟨font⟩
`\pdfobj` ⟨object type spec⟩ (h, v, m)
`\pdfoutline` [⟨attr spec⟩] ⟨action spec⟩ [`count` ⟨integer⟩] ⟨general text⟩ (h, v, m)
`\pdfprimitive` ⟨control sequence⟩
`\pdfrefobj` ⟨object number⟩ (h, v, m)
`\pdfrefxform` ⟨object number⟩ (h, v, m)
`\pdfrefximage` ⟨object number⟩

\pdfresettimer

\pdfrestore

\pdfrunninglinkoff

\pdfrunninglinkon

\pdfsave

\pdfsavepos (h, v, m)

\pdfsetmatrix

\pdfsetrandomseed ⟨number⟩

\pdfspacefont ⟨general text⟩

\pdfstartlink [⟨rule spec⟩] [⟨attr spec⟩] ⟨action spec⟩ (h, m)

\pdfthread [⟨rule spec⟩] [⟨attr spec⟩] ⟨id spec⟩ (h, v, m)

\pdftrailer ⟨general text⟩

\pdftrailerid ⟨general text⟩

\pdftstartthread [⟨rule spec⟩] [⟨attr spec⟩] ⟨id spec⟩ (v, m)

\pdfxform [⟨attr spec⟩] [⟨resources spec⟩] ⟨box number⟩ (h, v, m)

\pdfximage [⟨rule spec⟩] [⟨attr spec⟩] [⟨page spec⟩] [⟨pdf box spec⟩] [⟨colorspace spec⟩]
       ⟨general text⟩ (h, v, m)

\quitvmode

\special {pdf: ⟨text⟩}

\special {pdf:direct: ⟨text⟩}

\special {pdf:page: ⟨text⟩}

\special [shipout] {⟨text⟩}

\vadjust [⟨pre spec⟩] ⟨filler⟩ {⟨vertical mode material⟩ } (h, m)

B.7   General definitions and syntax rules

⟨general text⟩ → {⟨balanced text⟩ }

⟨attr spec⟩ → attr ⟨general text⟩

⟨resources spec⟩ → resources ⟨general text⟩

⟨rule spec⟩ → (width | height | depth) ⟨dimen⟩ [⟨rule spec⟩]

⟨object type spec⟩ → reserveobjnum |
       [useobjnum ⟨number⟩]
       [stream [⟨attr spec⟩]] ⟨object contents⟩

⟨annot type spec⟩ → reserveobjnum |
       [useobjnum ⟨number⟩] [⟨rule spec⟩] ⟨general text⟩

⟨object contents⟩ → ⟨file spec⟩ | ⟨general text⟩

⟨xform attr spec⟩ → [⟨attr spec⟩] [⟨resources spec⟩]

⟨image attr spec⟩ → [⟨rule spec⟩] [⟨attr spec⟩] [⟨page spec⟩] [⟨colorspace spec⟩]
       [⟨pdf box spec⟩]

⟨outline spec⟩ → [⟨attr spec⟩] ⟨action spec⟩ [count ⟨number⟩] ⟨general text⟩

⟨action spec⟩ → user ⟨user-action spec⟩ | goto ⟨goto-action spec⟩ |
       thread ⟨thread-action spec⟩

⟨user-action spec⟩ → ⟨general text⟩

⟨goto-action struct spec⟩ → struct (⟨numid⟩ | nameid | ⟨general text⟩)

⟨goto-action spec⟩ → [⟨goto-action struct spec⟩] ⟨numid⟩ |
       [⟨file spec⟩] [⟨goto-action struct spec⟩] ⟨nameid⟩ |

[⟨file spec⟩] [⟨goto-action struct spec⟩] [⟨page spec⟩] ⟨general text⟩ |
⟨file spec⟩ [⟨goto-action struct spec⟩] ⟨nameid⟩ ⟨newwindow spec⟩ |
⟨file spec⟩ [⟨goto-action struct spec⟩] [⟨page spec⟩] ⟨general text⟩ ⟨newwindow spec⟩

⟨thread-action spec⟩ → [⟨file spec⟩] ⟨numid⟩ | [⟨file spec⟩] ⟨nameid⟩

⟨colorspace spec⟩ → colorspace ⟨number⟩

⟨pdf box spec⟩ → mediabox | cropbox | bleedbox | trimbox | artbox

⟨map spec⟩ → {[⟨map modifier⟩] ⟨balanced text⟩ }

⟨map modifier⟩ → + | = | -

⟨numid⟩ → num ⟨number⟩

⟨nameid⟩ → name ⟨general text⟩

⟨newwindow spec⟩ → newwindow | nonewwindow

⟨dest spec⟩ → [struct ⟨number⟩] (⟨numid⟩ | ⟨nameid⟩) ⟨dest type⟩

⟨dest type⟩ → xyz [zoom ⟨number⟩] | fitr ⟨rule spec⟩ |
        fitbh | fitbv | fitb | fith | fitv | fit

⟨thread spec⟩ → [⟨rule spec⟩] [⟨attr spec⟩] ⟨id spec⟩

⟨id spec⟩ → ⟨numid⟩ | ⟨nameid⟩

⟨file spec⟩ → file ⟨general text⟩

⟨page spec⟩ → page ⟨number⟩

⟨expand spec⟩ → ⟨stretch⟩ ⟨shrink⟩ ⟨step⟩ [autoexpand]

⟨stretch⟩ → ⟨number⟩

⟨shrink⟩ → ⟨number⟩

⟨step⟩ → ⟨number⟩

⟨pre spec⟩ → pre

⟨stack action⟩ → set | push | pop | current

A ⟨general text⟩ is expanded immediately, like \special in traditional TEX, unless explicitly mentioned otherwise.

Some of the object and image-related primitives can be prefixed by \immediate.

Appendix C

Abbreviations used in this manual

In this document we use numerous abbreviations. For convenience we give their meanings here.

AFM	Adobe Font Metrics
ASCII	American Standard Code for Information Interchange
ConTeXt	general purpose macro package
CTAN	global TeX archive server
DVI	native TeX DeVice Independent file format
encTeX	encTeX extension to TeX
epstopdf	EPS to PDF conversion tool
EPS	Encapsulated PostScript
Eplain	Expanded plain TeX format
ε-TeX	a stable extension to TeX
EXIF	Exchangeable Image File format (JPEG file variant)
Ghostscript	PostScript and PDF language interpreter
GNU	GNU's Not Unix
HZ	Hermann Zapf's paragraph-breaking optimizations
ISO	International Organization for Standardization
JBIG2	Joint Bi-level Image Experts Group image format, version 2
JBIG	Joint Bi-level Image Experts Group image format
JFIF	JPEG File Interchange Format
JPEG	Joint Photographic Experts Group
LaTeX	general-purpose macro package
LMTX	the LuaMetaTeX engine
LuaTeX	the LuaTeX engine
MacTeX	TeX Live on the Mac
METAFONT	graphic programming environment, bitmap output
MetaPost	graphic programming environment, vector output
MiKTeX	TeX distribution for Windows
mlTeX	MLTeX extension to TeX
mptopdf	MetaPost to PDF conversion tool
PDF/A	PDF A/* standards
pdfeTeX	ε-TeX extension supporting PDF output
pdfLaTeX	LaTeX format using pdfTeX, producing PDF
pdfTeX	TeX extension supporting PDF output
PDF	Portable Document Format
Perl	Perl programming environment

PFA	Adobe PostScript Font format, ASCII
PFB	Adobe PostScript Font format, binary
PK	Packed bitmap font
PNG	Portable Network Graphics
POSIX	Portable Operating System Interface
PostScript	general graphics language
PStoPDF	PostScript to PDF converter (on top of Ghostscript)
RGB	red–green–blue color specification
TCX	TeX Character Translation
TDS	TeX Directory Standard
Texinfo	GNU documentation format
TeX Live	TeX Live distribution (cross-platform)
TeX	typographic language and program
TFM	TeX Font Metrics
TIFF	Tagged Interchange File format
TUG	TeX Users Group, `tug.org`
Unix	Unix platform
url	Uniform Resource Locator
UTF-8	Uniform Resource Locator
Web2C	Implementation framework for TeX and friends
WEB	literate programming environment
Windows	Microsoft Windows platform
XeTeX	the XeTeX engine

Appendix D

GNU Free Documentation License (v1.2)

Version 1.2, November 2002
Copyright © 2000,2001,2002 Free Software Foundation, Inc.

51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondarily, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "**Document**", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "**you**". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "**Modified Version**" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "**Secondary Section**" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "**Invariant Sections**" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "**Cover Texts**" are certain short passages

of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "**Transparent**" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "**Opaque**".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "**Title Page**" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "**Entitled XYZ**" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "**Acknowledgements**", "**Dedications**", "**Endorsements**", or "**History**".) To "**Preserve the Title**" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location

from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.

B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.

C. State on the Title page the name of the publisher of the Modified Version, as the publisher.

D. Preserve all the copyright notices of the Document.

E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.

F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.

G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.

H. Include an unaltered copy of this License.

I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.

J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.

K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.

L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.

M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.

N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.

O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied

from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties–for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but

you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See https://www.gnu.org/licenses/.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

ADDENDUM: How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

> Copyright © YEAR YOUR NAME. Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled "GNU Free Documentation License".

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the "with . . . Texts." line with this:

> with the Invariant Sections being LIST THEIR TITLES, with the Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.