

PDFTEX

users manual

Hàn Thê Thành Sebastian Rahtz Hans Hagen Hartmut Henkel

Contents

1	Introduction	9	Character translation
2	About PDF	10	Limitations of PDF \TeX
3	Getting started		
4	Macro packages supporting PDF \TeX		Abbreviations
5	Setting up fonts		Examples of HZ and protruding
6	Formal syntax specification		Additional PDF keys
7	New primitives		Colophon
8	Graphics and color		GNU Free Documentation License

1 Introduction

The main purpose of the pdf \TeX project is to create and maintain an extension of \TeX that can produce pdf directly from \TeX source files and improve/enhance the result of \TeX typesetting with the help of pdf. When pdf output is not selected, pdf \TeX produces normal dvi output, otherwise it generates pdf output that looks identical to the dvi output. An important aspect of this project is to investigate alternative justification algorithms (e. g. a font expansion algorithm akin to the hz micro–typography algorithm by Prof. Hermann Zapf), optionally making use of multiple master fonts.

pdf \TeX is based on the original \TeX sources and Web2c, and has been successfully compiled on Unix, Win32 and MSDos systems. It is under active development, with new features trickling in. Great care is taken to keep new pdf \TeX versions backward compatible with earlier ones.

For some years there has been a ‘moderate’ successor to \TeX available, called ε - \TeX . Because mainstream macro packages such as L^A \TeX have started supporting this welcome extension, pdf \TeX also is available as pdf ε \TeX . Although in this document we will speak of pdf \TeX , we advise users to use pdf ε \TeX when available. That way they get the best of all worlds and are ready for the future. Starting with \TeX Live 2004, that future has arrived: pdf ε \TeX is now the primary \TeX engine.

Other extensions to pdf \TeX are ML \TeX and enc \TeX ; recent pdf \TeX engines have these often included.

pdf \TeX is maintained by Hàn Thê Thành, Martin Schröder, Hartmut Henkel, Hans Hagen and others. The pdf \TeX homepage is [http://www.pdf \$\varepsilon\$ tex.org](http://www.pdfεtex.org). Please send pdf \TeX comments and bug reports to the mailing list [pdf \$\varepsilon\$ tex@tug.org](mailto:pdfεtex@tug.org).

We thank all readers who send us corrections and suggestions. We also wish to express the hope that pdfT_EX will be of as much use to you as it is to us. Since pdfT_EX is still being improved and extended, we suggest you to keep track of updates.

1.1 About this manual

This manual revision tries to keep track with the recent pdfT_EX development up to version 1.20b. Main text updates were done regarding the new configuration scheme, font mapping, and new or updated primitives. The primary repository for the manual and its sources is at <http://sarovar.org/projects/pdftex/>. Copies in pdf format can also be found at the CTAN network in directory `ctan:systems/pdftex`.

Thanks to Karl Berry for proof reading and submitting a long changes list. New errors might have slipped in afterwards by the editor. Please send questions or suggestions by email to pdftex@tug.org.

1.2 Legal Notice

Copyright © 1996–2005 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”.

2 About PDF

The cover of this manual lists an almost minimal pdf file generated by pdfT_EX, with the corresponding source file on the next page. Unless compression is enabled, such a pdf file is rather verbose and readable. The first line specifies the version used; currently pdfT_EX produces level 1.4 output. Viewers are supposed to silently skip over all elements they cannot handle.

A pdf file consist of objects. These objects can be recognized by their number and keywords:

```
8 0 obj << /Type /Catalog /Pages 6 0 R >> endobj
```

Here `8 0 obj ... endobj` is the object capsule. The first number is the object number. Later we will see that pdfTeX gives access to this number. One can for instance create an object by using `\pdfobj` after which `\pdflastobj` returns the number. So

```
\pdfobj{/Type /Catalog /Pages 6 0 R}
```

inserts an object into the file, while `\pdflastobj` returns the number pdfTeX assigned to this object. The sequence `6 0 R` is an object reference, a pointer to another object (no. 6). 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. The version numbers permit a roll-back. (An example of a graphic editor that uses pdf as storage format is Adobe's Illustrator.)

In general this rather direct way of pushing objects in the files 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. Unless such direct objects are part of something larger, they will end up as isolated entities, not doing any harm but not doing any good either.

When a viewer opens a pdf file, it first goes to the end of the file. There it finds the keyword `startxref`, the signal where to look for the so called 'object cross reference table'. This table provides fast access to the objects that make up the file. The actual starting point of the file is defined after the `trailer`. The `/Root` entry points to the catalog. In this catalog the viewer can find the page list. In our example we have only one page. The trailer also holds an `/Info` entry, which tells a bit more about the document. Just follow the thread:

```
/Root → object 8 → /Pages → object 6 → /Kids → object 2 → page content
```

As soon as we add annotations, a fancy word for hyperlinks and the like, some more entries are 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`. Let's take a closer look at this stream. First 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 switch can be recognized as `/F...`. The actual text goes between `()` so that it creates a PostScript string. 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, the words comes out rather fragmented, mainly because a lot of kerning takes place. 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.

This one page example uses an Adobe Times–Roman font. This is one of the 14 so–called standard fonts that are always present in the viewer application, and therefore need not be embedded in the pdf file. However, when we use for instance Computer Modern Roman, we have to make sure that this font is available, and the best way to do this is to embed it. Just let your eyes follow the object thread and see how a font is described. The only thing removed from this example is the (partially) embedded glyph description file, which for the 14 standard fonts is not needed.

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 (`/Type /Page`) shows that a mediabox is part of the page description. A mediabox acts like the bounding box in a PostScript file. pdfTeX users have access to this object by `\pdfpageattr`.

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. Those who, after this introduction, feel unsure how to proceed, are advised to read on but skip section 7. Before we come to that section, we will describe how to get started with pdfTeX.

3 Getting started

This section describes the steps needed to get pdfTeX running on a system where pdfTeX is not yet installed. Nowadays virtually all TeX distributions have pdfTeX as a component, such as TeX Live, teTeX, fpTeX, MikTeX, and CMacTeX. The ready to run TeX Live distribution comes with pdfTeX versions for many Unix Win32, and Mac OS X systems; more information can be found at <http://www.tug.org/tex-live/>. teTeX by Thomas Esser is a source distribution with an automated compilation process for Unix systems; see <http://www.tug.org/teTeX/>. For Win32 systems there are also two separate distributions that contain pdfTeX, both in `ctan:systems/win32`: fpTeX by Fabrice Popineau and MikTeX by Christian Schenk. So when you use any of these distributions, you don't need to bother with the pdfTeX installation procedure in the next sections.

If there is no precompiled binary of pdfTeX 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 T_EX system, e. g. t_EX, into which the freshly compiled pdfT_EX will be integrated. Note that the installation description in this manual is Web2c-specific.

3.1 Getting sources and binaries

The latest sources of pdfT_EX are currently distributed for compilation on Unix systems (including Linux), and Win32 systems (Windows 95, 98, NT, 2000, XP). The primary location where one can fetch the latest released code is at the developers' homepage <http://sarovar.org/projects/pdftex/>, where you also find bug tracking information, and the manual sources. Download the pdfT_EX tarball from the directory:

```
http://pdftex.sarovar.org/current/
```

The pdfT_EX sources can also be found at their canonical place in the CTAN network, `ctan:systems/pdftex`. Separate pdfT_EX binaries for various systems might also be available, check out the subdirectories below `ctan:systems`.

3.2 Compiling

The compilation is expected to be easy on Unix-like systems and can be described best by example. Assuming that the file `pdftex.tar.bz2` is downloaded to some working directory, e. g. `$HOME/pdftex`, on a Unix system the following steps are needed to compile pdfT_EX:

```
cd $HOME/pdftex
bzip2 -d pdftex-1.20b.tar.bz2 | tar xvf -
cd pdftex-1.20b
./Build
```

The binaries `pdftex` and `pdfetex` are then built in the subdirectory `build/texk/web2c`.

Apart from the binaries pdfT_EX and pdf_ET_EX the compilation also produces a few other files which are needed for running both pdfT_EX versions:

`pdftex.pool`, `pdfetex.pool` The pool files, needed for creating formats, located in `build/texk/web2c`

3.3 Placing files

The next step is to put the freshly compiled binaries and pool files into their proper places within the tds structure of the T_EX system. Put the files `pdftex` and `pdfetex` into the directory (e. g. for a typical t_EX system) `/usr/local/tEX/bin/i686-pc-linux-gnu`, and the pool files into `/usr/local/tEX/share/texmf/web2`.

Don't forget to do a `texconfig init` afterwards, so that all formats are regenerated with the fresh binaries.

3.4 Setting search paths

Web2c-based programs, including pdfT_EX, use the Web2c run-time configuration file called `texmf.cnf`. The location of this file is the appropriate position within the tds tree relative to the place of the pdfT_EX binary; on a t_EX system, file `texmf.cnf` typically is located either in directory `texmf/web2c` or `texmf-local/web2c`. The path to file `texmf.cnf` can also be set up by the environment variable `TEXMFCNF`.

Next you might need to edit `texmf.cnf` so that pdfT_EX can find all necessary files, but the `texmf.cnf` files coming with the major T_EX distributions should already be set up for normal use. You might check into the file `texmf.cnf` to see where the various bits and pieces are going.

pdfT_EX uses the search path variables shown in [table 1](#).

<code>TEXMFOUTPUT</code>	Normally, pdfT _E X puts its output files in the current directory. If any output file cannot be opened there, it tries to open it in the directory specified in the environment variable <code>TEXMFOUTPUT</code> . There is no default value for that variable. For example, if you type <code>pdfetex paper</code> and the current directory is not writable, if <code>TEXMFOUTPUT</code> has the value <code>/tmp</code> , pdfT _E X attempts to create <code>/tmp/paper.log</code> (and <code>/tmp/paper.pdf</code> , if any output is produced.)
<code>TEXINPUTS</code>	This variable specifies where pdfT _E X finds its input files. Image files are considered input files and searched for along this path.
<code>TEXFORMATS</code>	Search path for format (<code>.fmt</code>) files.
<code>TEXPOOL</code>	Search path for pool (<code>.pool</code>) files.
<code>ENCFONTS</code>	Search path for encoding (<code>.enc</code>) files.

used for	texmf.cnf
output files	TEXMFOUTPUT
input files, images	TEXINPUTS
format files	TEXFORMATS
text pool files	TEXPOOL
encoding files	ENCFONTS
font map files	TEXFONTPMAPS
tfm files	TFMFONTS
virtual fonts	VFFONTS
Type 1 fonts	T1FONTS
TrueType fonts	TTFONTS
pixel fonts	PKFONTS

Table 1 The Web2c variables.

TEXFONTPMAPS	Search path for font map (<code>.map</code>) files.
TFMFONTS	Search path for font metric (<code>.tfm</code>) files.
VFFONTS	Search path for virtual font (<code>.vf</code>) files. Virtual fonts are fonts made up of other fonts. Because pdfTeX produces the final output code, it must consult those files.
T1FONTS	Search path for Type 1 font files (<code>.pfa</code> and <code>.pfb</code>). These outline (vector) fonts are to be preferred over bitmap pk fonts. In most cases Type 1 fonts are used and this variable tells pdfTeX where to find them.
TTFONTS	Search path for TrueType font (<code>.ttf</code>) files. Like Type 1 fonts, TrueType fonts are also outlines.
PKFONTS	Search path for packed (bitmap) font (<code>.pk</code>) files. Unfortunately bitmap fonts are still displayed poorly by some pdf viewers, so when possible one should use outline fonts. When no outline is available, pdfTeX tries to locate a suitable pk font (or invoke a process that generates it).

3.5 The PDF_TE_X configuration

One has to keep in mind that, as opposed to T_EX with its dvi output, the pdfT_EX program does not require a separate postprocessing stage to transform the T_EX input into a pdf file. As a consequence, all data needed for building a ready pdf page must be available during the pdfT_EX run, in particular information on media dimensions and offsets, graphics files for embedding, and font information (font files, encodings).

When T_EX builds a page, it places items relative to 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 pdfT_EX therefore requires that they be defined at the beginning of the pdfT_EX run. As with pages described by PostScript, the pdf reference point is in the lower-left corner.

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

A list of pdfT_EX primitives relevant for setting up the pdfT_EX engine is given in [table 2](#). These are described in detail within later sections. [Figure 1](#) shows a recent configuration file (`pdfTEXconfig.tex`) in T_EX format, using the primitives from [table 2](#), which typically is read in during the format building process. It enables pdf output, sets a high compression level for reducing pdf file size, sets item placement precision, paper dimensions, and offsets, sets the default pixel density for pk font inclusion, and sets the pdf version number to appear in the pdf file. The default values are chosen so that pdfT_EX often can be used (e. g. in `-ini` mode) even without setting any parameters.

Independent of whether such a configuration file is read or not, the first action in a pdfT_EX run is that the program reads 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. pool and hash size), and other general parameters.

internal name	type	default	comment
<code>\pdfoutput</code>	integer	0	dvi
<code>\pdfadjustspacing</code>	integer	0	off
<code>\pdfcompresslevel</code>	integer	9	best
<code>\pdfdecimaldigits</code>	integer	4	max.
<code>\pdfmovechars</code>	integer	0	off
<code>\pdfimageresolution</code>	integer	72	dpi
<code>\pdfpkresolution</code>	integer	0	72 dpi
<code>\pdfuniqueresname</code>	integer	0	
<code>\pdfprotrudechars</code>	integer	0	
<code>\pdfoptionpdfminorversion</code>	integer	4	pdf 1.4
<code>\pdfoptionalwaysusepdfpagebox</code>	integer	0	
<code>\pdfoptionpdfinclusionerrorlevel</code>	integer	0	
<code>\pdfhorigin</code>	dimension	1 in	
<code>\pdfvorigin</code>	dimension	1 in	
<code>\pdfpagewidth</code>	dimension	0 pt	
<code>\pdfpageheight</code>	dimension	0 pt	
<code>\pdflinkmargin</code>	dimension	0 pt	
<code>\pdfdestmargin</code>	dimension	0 pt	
<code>\pdfthreadmargin</code>	dimension	0 pt	
<code>\pdfmapfile</code>	text	<code>pdftex.map</code>	

Table 2 The set of pdfTeX configuration parameters.

3.6 Creating format files

Both pdfTeX and pdf_εTeX engines allow building formats for dvi and pdf output in the same way as the classical TeX engine does for dvi. Format generation 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 then inherits this setting, so that a later call to pdfTeX with this format starts in the preselected mode (which still can be

```

% Set pdfTeX parameters for pdf mode
% (replacing pdftex.cfg file).
% Thomas Esser, 2004. public domain.
\pdfoutput=1
\pdfcompresslevel=9
\pdfdecimaldigits=3
\pdfpagewidth=210 true mm
\pdfpageheight=297 true mm
\pdfhorigin=1 true in
\pdfvorigin=1 true in
\pdfpkresolution=1200
\pdfoptionpdfminorversion=4
\endinput

```

Figure 1 A typical configuration file (pdftexconfig.tex).

```

% Thomas Esser, 1998, 2004. public domain.
\ifx\pdfoutput\undefined
\else
  \ifx\pdfoutput\relax
  \else
    \input pdftexconfig
    \pdfoutput=0
  \fi
\fi
\input etex.src
\dump
\endinput

```

Figure 2 File etex.ini for ϵ -TeX format with dvi output.

overrun then). A format file can be read in only by the engine that has generated it; a format incompatible with an engine leads

```
\ifx\pdfoutput\undefined
\else
\ifx\pdfoutput\relax
\else
\input pdftexconfig
\pdfoutput=1
\fi
\fi
\scrollmode
\input latex.ltx
\endinput
```

Figure 3 File `pdflatex.ini` for L^AT_EX format with pdf output.

to a fatal error. Often the pdfT_EX program is a mere link to the pdf_ET_EX engine; then also a pdfT_EX call generates an extended format.

It is customary to package the configuration and macro file input into a `.ini` file. E. g., the file `etex.ini` in [figure 2](#) is for generating an ϵ -T_EX format with dvi output (it contains a few comparisons to be safe also for T_EX engines). A similar file `pdflatex.ini` can be used for generating a L^AT_EX format with pdf output; refer to [figure 3](#). One can see how the primitive `\pdfoutput` is used to override the output mode set by file `pdftexconfig.tex`. The corresponding pdfT_EX calls for format generation are:

```
pdfetex -ini *etex.ini
pdfTeX -ini pdflatex.ini
```

These calls produce format files `etex.fmt` and `pdflatex.fmt`, as the default format file name is taken from the input file name. You can overrule this with the `-jobname` option. The asterisk (*) in the first example line tells the pdf_ET_EX engine to go into extended `-ini` mode. Otherwise, it stays in non-extended mode. So, if you want a pdfL^AT_EX format with pdf output and ϵ -T_EX extensions available (format file `pdfelatex.fmt`), you would need to type e. g.:

```
pdfetex -ini -jobname=pdfelatex *pdflatex.ini
```

In ConT_EXt the generation depends on the interface used. A format using the English user interface is generated with

```
pdfetex -ini cont-en
```

When properly set up, one can also use the ConT_EXt command line interface T_EXexec to generate one or more formats, like:

```
texexec --make en
```

for an English format, or

```
texexec --make en de
```

for an English and German one. Most users will simply say:

```
texexec --make --all [--alone]
```

and so generate the T_EX and METAPOST related formats that ConT_EXt needs. Whatever macro package used, the formats should be placed in the TEXFORMATS path.

3.7 Testing the installation

When everything is set up, you can test the installation. In the distribution there is a plain T_EX test file `example.tex`. Process this file by typing:

```
pdftex example
```

If the installation is ok, this run should produce a file called `example.pdf`. The file `example.tex` is also a good place to look for how to use pdfT_EX's new primitives.

3.8 Common problems

The most common problem with installations is that pdfT_EX complains that something cannot be found. In such cases make sure that TEXMFCNF is set correctly, so pdfT_EX can find `texmf.cnf`. The next best place to look/edit is the file `texmf.cnf`. When still in deep trouble, set `KPATHSEA^DEBUG=255` before running pdfT_EX or run pdfT_EX with option `-k 255`. This will cause pdfT_EX to write a lot of debugging information that can be useful to trace problems. 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 read `pdftex.pool`; bad path?

`TEXMFCNF` is not set correctly and so `pdfTeX` cannot find `texmf.cnf`, or `TEXPOOL` in `texmf.cnf` doesn't contain a path to the pool file `pdftex.pool` or `pdfetex.pool` when you use `pdfTeX`.

- You have to increase `POOLSIZE`.

`pdfTeX` cannot find `texmf.cnf`, or the value of `pool-size` specified in `texmf.cnf` is not large enough and must be increased. If `pool-size` is not specified in `texmf.cnf` then you can add something like

```
pool-size = 500000
```

- 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`.

- ---! `xx.fmt` was written by `tex`
Fatal format file error; I'm stymied

This appears e. g. if you forgot to regenerate the `.fmt` files after installing a new version of the `pdfTeX` binary and `pdftex.pool`. The first line tells by which engine the offending format was generated.

- `TEX.POOL` doesn't match; TANGLE me again!
`TEX.POOL` doesn't match; TANGLE me again (or fix the path).

This might appear if you forgot to install the proper `pdftex.pool` when installing a new version of the `pdfTeX` binary. E. g. under `teTeX` then run `texconfig init`.

- `pdfTeX` cannot find one or more map files (`*.map`), encoding vectors (`*.enc`), virtual fonts, Type 1 fonts, TrueType 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. See above which variables pdfTeX needs apart from the ones TeX uses.

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.

Normally the page content takes one object. This means that one seldom finds more than a few hundred objects in a simple file. This document for instance uses about 650 objects. In demanding applications this number can grow quite rapidly, especially when one uses a lot of widget annotations, shared annotations or other shared things. In these situations in `texmf.cnf` one can enlarge pdfTeX's internal object table by adding a line in `texmf.cnf`, for instance:

```
obj-tab-size = 400000
```

4 Macro packages supporting PDFTeX

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. Typical tasks are handling color, graphics, hyperlink support, threading, font-inclusion, as well as page imposition and manipulation. All these pdf-specific tasks can be commanded by pdfTeX's own primitives (a few also by a pdfTeX-specific `\special{pdf: ...}` primitive). 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 only for non-empty `\special{}` commands.

When a macro package already written for classical TeX with dvi output is to be modified for use with pdfTeX, it is very helpful to get some insight to what extent pdfTeX-specific support is needed. This info can be gathered e. g. by outputting the various `\special` commands as `\message`. Simply type

```
\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.

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 L^AT_EX users, Sebastian Rahtz's and Heiko Oberdiek's `hyperref` package 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 page size is set up correctly. Bookmarks are created to match the table of contents.
- The standard L^AT_EX `graphics`, `graphicx`, and `color` packages also have automatic pdfTeX support, which allow use of color, text rotation, and graphics inclusion commands.
- The ConT_EXt macro package by Hans Hagen has very full support for pdfTeX in its generalized hypertext features. Support for pdfTeX is implemented as a special driver, and is invoked by typing `\setupoutput [pdf tex]` or feeding T_EX_{exec} with the `--pdf` option.
- pdf from Texinfo documents can be created by running pdfTeX on the Texinfo file, instead of T_EX. Alternatively, run the shell command `texi2pdf` instead of `texi2dvi`.
- A small modification of `webmac.tex`, called `pdfwebmac.tex`, allows production of hyperlinked pdf versions of the program code written in web.

Some nice samples of pdfTeX output can be found at <http://www.pdfTeX.org> and <http://www.pragma-ade.com>.

5 Setting up fonts

pdfTeX can work with Type 1 and TrueType fonts, but a source must be available for all fonts used in the document, except for the 14 standard fonts supplied by the pdf reader (Times, Helvetica, Courier, Symbol and Dingbats). 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 very poorly in (older versions of) Acrobat Reader. Given the free availability of Type 1 versions of all the Computer Modern fonts, and the ability to use standard PostScript fonts, there is rarely a need to use bitmap fonts in pdfTeX.

5.1 Map files

Font map files provide the connection between T_EX tfm font files and the outline font file names. They contain also information about re-encoding arrays, partial downloading, and character transformation parameters (like `SlantFont` and `ExtendFont`). Those map files were first created for dvi postprocessors. But, as pdfT_EX 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 pdfT_EX is in dvi mode. Pixel fonts can be used without being listed in the map file.

By default, pdfT_EX reads the map file `pdf tex .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 an individual line. The syntax of each line is upward-compatible with `dvips` map files and can contain the following (some are optional) fields: *tfmname*, *basename*, *fontflags*, *special*, *encodingfile*, and *fontfile*; explanations follow.

It is mandatory that *tfmname* is the first field. If a *basename* is given, it must be the second field. Similarly if *fontflags* is given it must be the third field (if *basename* is present) or the second field (if *basename* is left out). It is possible to mix the positions of *special*, *encodingfile*, and *fontfile*, however the first three fields must be given in fixed order.

tfmname sets the name of the tfm file for a font – the name T_EX sees. This name must always be given.

basename sets the base (PostScript) font name. The *basename* field is checked against the `BaseName` entry of fonts coming with embedded pdf files. If there is a match, the font will be removed from the embedded file, and a local font is opened, which will contain the glyphs from the embedded file. This collecting mechanism helps keeping the resulting pdf file size small, if many files with similar fonts are to be embedded. Therefore it is recommended always to set the *basename* field.

If a *basename* field is given, also a *fontfile* field must be there, unless the *basename* matches one of the 14 standard font names; then the *fontfile* field is optional. If the *fontfile* name is given, this font will be embedded (depending on flags, see below). If the *fontfile* name for a standard font is missing, the font will be quietly left out, which is fine, as pdf viewers will later render the text with their own versions of the font.

fontflags specify some 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 can adapt their rendering to these flags, especially when they substitute a replacements for not embedded fonts.

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. Table 3 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-cap font
18	Small-cap font
19	Force bold at small text sizes
20–32	Reserved

Table 3 The meaning of flags in the font descriptor.

All characters in a *fixed-width* font 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-cap* font, which is typically used for display purposes such as titles or headlines, contains no lowercase letters. It differs from a *small-cap* 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 font flags are not given, pdfTEX treats it as being 4, a symbolic 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 troubles 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).

special instructions can be used to manipulate fonts similar to the way dvips does. Currently only the keywords `SlantFont` and `ExtendFont` are interpreted, other instructions (as `ReEncodeFont` with parameters, see *encoding* below) are just ignored. The permitted `SlantFont` range is $-1..1$; for `ExtendFont` it's $-2..2$. The block of *special* instruction must be enclosed by double quotes `""`.

encoding specifies the name of the file containing the external encoding vector to be used for the font. The file name may be preceded by a `<`, but the effect is the same. 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. It may be omitted if you are sure that the font resource has the correct built-in encoding. In general this option is highly preferred and is *required* when subsetting a TrueType font.

fontfile sets the name of the font source file. This must be a Type 1 or TrueType font file. The font file name can be preceded by one or two special characters, which says how the font file should be handled.

- If the font file name is preceded by a `<` the font file will be partially downloaded, meaning that only used glyphs (characters) are embedded to the font. 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. Partial 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 the ones that are not used in the document. Apart from causing large size pdf output, this option may cause troubles with TrueType fonts, so it is not recommended. It might be useful in case the font is atypical and can not be subsetted well by pdfTEX. *Beware: some font vendors forbid full font inclusion.*

- If nothing precedes the font file name, the font file is read but nothing is embedded, only the font parameters are extracted to generate the so-called font descriptor, which is used by the pdf reader to simulate the font if needed. This option is useful only when you do not want to embed the font (i. e. to reduce the output size), but wish to use the font metrics and let the pdf reader 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, which is not very useful.
- If the font file name is preceded by a !, the font is not read at all, and is assumed to be available on the system. This option can be used to create pdf files which do not contain embedded fonts. The pdf output then works only on systems where the resource of the used font is available. It's not very useful for document exchange, as the pdf is not 'portable' at all. On the other hand it is very useful when you wish to speed up running of pdfT_EX during interactive work, and only in a final version embed all used fonts. Don't over-estimate gain in speed and when distributing files, always embed the fonts! This feature requires the pdf reader to have access to installed fonts on the system. This has been tested on Win95 and Unix (Solaris).

When one suffers from invalid lookups, for instance when pdfT_EX tries to open a .pfa file instead of a .pfb one, one can add the suffix to the filename. In this respect, pdfT_EX completely relies on the `kpathsea` libraries.

If a used font is not present in the map files, first pdfT_EX will look for a source with suffix .pgc, which is a so-called pgc source (pdf Glyph Container)¹. If no pgc source is available, pdfT_EX will try to use pk fonts as dvi drivers do, creating pk fonts on-the-fly if needed.

Lines containing nothing apart from *tfmname* stand for scalable Type 3 fonts. For scalable fonts as Type 1, TrueType and scalable Type 3 font, all the fonts loaded from a tfm at various sizes will be included only once in the pdf output. Thus if a font, let's say `csr10`, is described in one of the map files, then it will be treated as scalable. As a result the font source for `csr10` will be included only once for `csr10`, `csr10 at 12pt` etc. So pdfT_EX tries to do its best to avoid multiple downloading of identical font sources. Thus vector pgc fonts should be specified as scalable Type 3 in map files like:

```
csr10
```

¹This is a text file containing a pdf Type 3 font, created by METAPOST using some utilities by Hans Hagen. In general pgc files can contain whatever allowed in pdf page description, which may be used to support fonts that are not available in METAFONT. pgc fonts are not widely useful, as vector Type 3 fonts are not displayed very well in older versions of Acrobat Reader, but may be more useful when better Type 3 font handling is more common.

It doesn't hurt much if a scalable Type 3 font is not given in map files, except that the font source will be downloaded multiple times for various sizes, which causes a much larger pdf output. On the other hand if a font in the map files is defined as scalable Type 3 font and its pgc source is not scalable or not available, pdfTeX will use pk fonts instead; the pdf output is still valid but some fonts may look ugly because of the scaled bitmap.

To summarize this rather confusing story, we include a some example lines. First we use two fonts from the 14 standard fonts with font-specific encoding, i. e. no external encoding is given. In the first line, the fontfile is missing, so viewers will use their own font. The ZapfDingbats font is taken from the given font file.

```
psyr Symbol
pzdr ZapfDingbats <pzdr.pfb
```

Similarly, two standard fonts with an external encoding. The < preceding the encoding file name may be left out.

```
ptmr8r Times-Roman <8r.enc
ptmri8r Times-Italic <8r.enc <ptmri8a.pfb
```

A SlantFont is specified similarly as for dvips. The SlantFont or ExtendFont entries work only with embedded font files.

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

Download a font entirely without re-encoding:

```
pgsr8r GillSans <<pgsr8a.pfb
```

Partially download a font without re-encoding:

```
pgsr8r GillSans <pgsr8a.pfb
```

Do not read the font at all – the font is supposed to be installed on the system:

```
pgsr8r GillSans !pgsr8a.pfb
```

Entirely download a font with re-encoding:

```
pgsr8r GillSans <8r.enc <<pgsr8a.pfb
```

Partially download a font with re-encoding:

```
pgsr8r GillSans <8r.enc <pgsr8a.pfb
```

Sometimes we do not want to include a font, but need to extract parameters from the font file and re-encode the font as well. This only works for fonts with Adobe Standard Encoding. The font flags specify how such a font looks like, so Acrobat Reader can generate similar instance if the font resource is not available on the target system.

```
pgsr8r GillSans 32 <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
```

5.2 TrueType fonts

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`). Usage of `ttf2afm` is simple:

```
ttf2afm -e <encoding vector> -o <afm outputfile> <ttf input file>
```

A TrueType file can be recognized by its suffix `ttf`. The optional *encoding* specifies the encoding, which is the same as the encoding vector used in map files for pdfTeX and dvips. If the encoding is not given, all the glyphs of the afm output will be mapped to `.notdef`. `ttf2afm` writes the output afm to standard output. If we need to know which glyphs are available in the font, we can run `ttf2afm` without encoding to get all glyph names. The resulting afm file can be used to generate a tfm one by applying `afm2tfm`.

To use a new TrueType font the minimal steps may look like below. We suppose that `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 few limitations with TrueType fonts in comparison with Type 1 fonts:

- a. The special effects `SlantFont/ExtendFont` cannot be used.
- b. To subset a TrueType font, the font must be specified as re-encoded, therefore an encoding vector must be given.
- c. TrueType fonts coming with embedded pdf files are kept untouched; they are not replaced by local ones.

6 Formal syntax specification

This section formally specifies the pdfT_EX specific extensions to the T_EX macro programming language. All primitives are prefixed by `pdf` except for `\efcode`, `\lpcode`, and `\rprcode`. The general definitions and syntax rules follow after the list of primitives.

```
\pdfoutput (integer)
\pdfadjustspacing (integer)
\pdfcompresslevel (integer)
\pdfdecimaldigits (integer)
\pdfmovechars (integer)
\pdfimageresolution (integer)
\pdfpkresolution (integer)
\pdfuniqueresname (integer)
\pdfprotrudechars (integer)
\pdfoptionpdfminorversion (integer)
\pdfoptionalwaysusepdfpagebox (integer)
\pdfoptionpdfinclusionerrorlevel (integer)
\pdfhorigin (dimen)
\pdfvorigin (dimen)
\pdfpagewidth (dimen)
\pdfpageheight (dimen)
\pdflinkmargin (dimen)
\pdfdestmargin (dimen)
\pdfthreadmargin (dimen)
```

`\pdfpagesattr` (tokens)
`\pdfpageattr` (tokens)
`\pdfpageresources` (tokens)
`\pdftexrevision` (expandable)
`\pdftexbanner` (expandable)
`\pdffontname` \langle font \rangle (expandable)
`\pdffontobjnum` \langle font \rangle (expandable)
`\pdfincludechars` \langle font \rangle \langle general text \rangle (expandable)
`\pdftexversion` (read-only integer)
`\pdflastobj` (read-only integer)
`\pdflastxform` (read-only integer)
`\pdflastximage` (read-only integer)
`\pdflastximagepages` (read-only integer)
`\pdflastannot` (read-only integer)
`\pdflastxpos` (read-only integer)
`\pdflasttypos` (read-only integer)
`\pdflastdemerits` (read-only integer)
`\pdflastvbreakpenalty` (read-only integer)
`\pdfliteral` [`direct`] \langle general text \rangle (h, v, m)
`\pdfobj` \langle object type spec \rangle (h, v, m)
`\pdfrefobj` \langle object number \rangle (h, v, m)
`\pdfxform` [\langle xform attr spec \rangle] \langle box number \rangle (h, v, m)
`\pdfrefxform` \langle object number \rangle (h, v, m)
`\pdfximage` [\langle image attr spec \rangle] \langle general text \rangle (h, v, m)
`\pdfrefximage` \langle object number \rangle (h, v, m)
`\pdfannot` [\langle rule spec \rangle] \langle general text \rangle (h, v, m)
`\pdfstartlink` [\langle rule spec \rangle] [\langle attr spec \rangle] \langle action spec \rangle (h, m)
`\pdfendlink` (h, m)
`\pdfoutline` \langle outline spec \rangle (h, v, m)

`\pdfdest` \langle dest spec \rangle (h, v, m)
`\pdfthread` \langle thread spec \rangle (h, v, m)
`\pdfstartthread` \langle thread spec \rangle (v, m)
`\pdfendthread` (v, m)
`\pdfsavepos` (h, v, m)
`\pdfinfo` \langle general text \rangle
`\pdfcatalog` \langle general text \rangle [\langle open-action spec \rangle]
`\pdfnames` \langle general text \rangle
`\pdfmapfile` \langle map spec \rangle
`\pdfmapline` \langle map spec \rangle
`\pdffontattr` \langle font \rangle \langle general text \rangle
`\pdftrailer` \langle general text \rangle
`\pdffontexpand` \langle font \rangle \langle expand spec \rangle
`\efcode` \langle font \rangle \langle 8-bit number \rangle (integer)
`\lpcode` \langle font \rangle \langle 8-bit number \rangle (integer)
`\rprcode` \langle font \rangle \langle 8-bit number \rangle (integer)
`\vadjust` \langle pre spec \rangle \langle filler \rangle { \langle vertical material \rangle } (h, m)
`\special` \langle pdfspecial spec \rangle

The general definitions and syntax rules are given below. \langle general text \rangle is expanded immediately, like `\special` in traditional T_EX, unless explicitly mentioned otherwise.

\langle general text \rangle \rightarrow { \langle balanced text \rangle }
 \langle attr spec \rangle \rightarrow **attr** \langle general text \rangle
 \langle rule spec \rangle \rightarrow (**width** | **height** | **depth**) \langle dimension \rangle [\langle rule spec \rangle]
 \langle object type spec \rangle \rightarrow **reserveobjnum** |
 [**useobjnum** \langle number \rangle]
 [\langle attr spec \rangle] [**stream** [\langle attr spec \rangle]] \langle object contents \rangle
 \langle object contents \rangle \rightarrow \langle file spec \rangle | \langle general text \rangle
 \langle xform attr spec \rangle \rightarrow [\langle attr spec \rangle] [\langle resources spec \rangle]

<resources spec> → **resources** <general text>
 <image attr spec> → [<rule spec>] [<attr spec>] [<page 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 spec> → <numid> |
 [<file spec>] <nameid> |
 [<file spec>] [<page spec>] <general text> |
 <file spec> <nameid> <newwindow spec> |
 <file spec> [<page spec>] <general text> <newwindow spec>
 <thread-action spec> → [<file spec>] <numid> | [<file spec>] <nameid>
 <open-action spec> → **openaction** <action spec>
 <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> → <numid> <dest type> | <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>`
`<pdfspecial spec> → { [<pdfspecial id> [<pdfspecial modifier>]] <balanced text> }`
`<pdfspecial id> → pdf: | PDF:`
`<pdfspecial modifier> → direct:`

Some of the object and image related primitives can be prefixed by `\immediate`. More about that in the next sections.

7 New primitives

Here follows a short description of the primitives added by pdfTeX to the original TeX engine (other extensions by MLTeX and encTeX are not listed). One way to learn more about how to use these new primitives is to have a look at the file `example.tex` in the pdfTeX distribution.

Note that if the output is dvi then the pdfTeX specific dimension parameters are not used at all. However some pdfTeX integer parameters can affect the dvi as well as pdf output (currently `\pdfoutput` and `\pdfadjustspacing`).

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.

7.1 Document setup

▶ `\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 (unless the commandline option `-output-format=pdf` is used); all other primitives are optional. This parameter cannot be specified *after* shipping out the first page. In other words, if we want pdf output, we have to set `\pdfoutput` before pdfTeX ships out the first page.

When pdfTeX starts complaining about specials, one can be rather sure that a macro package is not aware of the pdf mode. A simple way of making macros aware of pdfTeX in 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 L^AT_EX and plain T_EX, is a cleaner way of doing this. Historically, the simple test `\ifx\pdfoutput\undefined` was defined; but nowadays, the pdfTeX engine is used in distributions even for non-pdf formats (e. g. L^AT_EX), so `\pdfoutput` may be defined even when the output format is dvi.

▶ `\pdfcompresslevel` (integer)

This integer parameter specifies the level of stream compression (text, in-line graphics, and embedded png images; all done by the `zlib` library). 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. Setting `\pdfcompresslevel=0` is great for pdf stream debugging.

▶ `\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 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 by typing:
`\twodigitMPoutput.`

▶ `\pdfmovechars` (integer)

This parameter specifies whether pdfTeX should try to move characters in range 0..31 to higher slots. When set to 1, this feature affects only those fonts that have all character codes below 128; for instance, Computer Modern. When set to 2 or higher pdfTeX will try to move those characters to free slots in encoding array, even if the font contains characters with code greater

than or equal to 128. This parameter is read when pdfTeX writes a character of a font to the pdf output, at which moment it has to decide whether to move the character or not.

▶ `\pdfpkresolution` (integer)

This integer parameter specifies the default resolution of embedded pk fonts and is read when pdfTeX downloads a pk font during finishing the pdf output. As bitmap fonts are still rendered poorly by some pdf viewers, it is best to use Type 1 fonts when available.

▶ `\pdfhorigin` (dimension)

This parameter can be used to set the horizontal offset 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 standard purposes, this parameter should always be kept at 1 true inch. If you want to shift text on the page, use TeX's own `\horigin` 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 typing `\pdfhorigin=1 true in` after issuing the `\mag` command. 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.

▶ `\pdfvorigin` (dimension)

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).

▶ `\pdfpagewidth` (dimension)

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. After magnification has been changed by the `\mag` primitive, check that this parameter reflects the wished true page width.

If the value is not given, the page width is calculated as $w_{\text{box being shipped out}} + 2 \times (\text{horigin} + \text{\hoffset})$. When part of the page falls off the paper or screen, you can be rather sure that this parameter is set wrong.

▶ `\pdfpageheight` (dimension)

Similar to the previous item, this dimension parameter specifies the page height of the pdf output. If not given, the page height will be calculated analogously to the above. After magnification has been changed by the `\mag` primitive, check that this parameter reflects the wished true page height.

▶ `\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 `/MediaBox`, the rectangle specifying the natural size of the page, `/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
  /MediaBox [0 0 612 792] } % the media size of all pages (in bp)
```

▶ `\pdfpageattr` (tokens)

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

7.2 The document info and catalog

▶ `\pdfinfo` (general text)

This primitive allows the user to add information to the document info section; if this information is provided, it can be extracted by Acrobat Reader (version 3.1: menu option *Document Information, General*). 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 are optional. Possible keys are `/Author`, `/CreationDate` (defaults to current date including time zone info), `/ModDate`, `/Creator` (defaults to TeX), `/Producer` (defaults to pdfTeX–1.20b), `/Title`, `/Subject`, and `/Keywords`.

`/CreationDate` and `/ModDate` are expressed in the form `D:YYYYMMDDhhmmssTZ. .`, where `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. An example of this format is shown below. For details please refer to the pdf Reference.

Multiple appearances of `\pdfinfo` will be concatenated. If a key is given more than once, then the first appearance will be used. An example of the use of `\pdfinfo` is:

```
\pdfinfo
{ /Title      (example.pdf)
  /Creator    (TeX)
  /Producer   (pdfTeX 1.20b)
  /Author     (Tom and Jerry)
  /CreationDate (D:20050106154343+01'00')
  /ModDate    (D:20050106155343+01'00')
  /Subject    (Example)
  /Keywords   (mouse, cat) }
```

► `\pdfcatalog` `<general text>` [`<open-action spec>`]

Similar to the document info section is the document catalog, where keys are `/URI`, which provides the base url of the document, and `/PageMode`, which determines how Acrobat displays the document on startup. The possibilities for the latter are explained in Table 4:

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

Table 4 Supported `/PageMode` values.

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

The `<openaction>` is the action provided when opening the document and is specified in the same way as internal links, see section 7.8. Instead of using this method, one can also write the open action directly into the catalog.

▶ `\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, otherwise the document can be invalid.

▶ `\pdftrailer` `<general text>`

This command puts its argument text verbatim into the file trailer dictionary.

▶ `\pdfoptionpdfminorversion` `(integer)`

This primitive sets the pdf version of the generated file and the latest allowed pdf version of included pdfs. E. g., `\pdfoptionpdfminorversion=3` tells pdfTeX to set the pdf version to 1.3 and allows only included pdf files with versions numbers up to 1.3. The default for `\pdfoptionpdfminorversion` is 4 for pdf version 1.4. If specified, this primitive must appear before any data is to be written to the generated pdf file, so you should put it at the very start of your files.

7.3 Fonts

▶ `\pdffontexpand` `` `<stretch>` `<shrink>` `<step>` `[autoexpand]`

This extension to T_EX's font definitions controls a pdfT_EX automatism called *font expansion*. We describe this by an example:

```
\font\somelfont=somelfile at 10pt
\pdffontexpand\somelfont 30 20 10 autoexpand
\pdfadjustspacing=2
```

The `30 20 10` means this: “hey T_EX, when line breaking is going badly, you may stretch the glyphs in this font as much as 3 % or shrink them by 2 %”. Because pdfT_EX uses internal data structures with fixed widths, each additional width also means an additional font. For practical reasons pdfT_EX uses discrete steps, in this example, 1 %. This means that for font `somelfile` up to 6 differently scaled alternatives may be used. When no step is specified, 0.5 % steps are used.

Roughly spoken, the trick is as follows. Consider a text typeset in triple column mode. When TEX cannot break a line in the appropriate way, the unbreakable parts of the word will stick into the margin. When $\text{pdf}\text{T}\text{E}\text{X}$ notes this, it will try to scale (shrink) the glyphs in that line using fixed steps, until the line fits. When lines are too spacy, the opposite happens: $\text{pdf}\text{T}\text{E}\text{X}$ starts scaling (stretching) the glyphs until the white space gaps is acceptable. This glyph stretching and shrinking is called *font expansion*.

The additional expanded fonts get artificial names by adding the font expansion value to the base font name, e. g. `somefile+10` for 1% stretch or `somefile-15` for 1.5% shrink. If the `autoexpand` option is not given, tfm files with these names and appropriate dimensions must be available. So, each expanded variant of a font must have its own tfm file! Expanded tfm names like `somefile+10` must not be mentioned in the map file (but the tfm name of the base font without expansion must be there). When no tfm file can be found, $\text{pdf}\text{T}\text{E}\text{X}$ will try to generate it by executing the script `mktexmf`, where available and supported.

The font expansion is greatly simplified, if the `autoexpand` option is there. Then no expanded tfm file versions are needed; instead, $\text{pdf}\text{T}\text{E}\text{X}$ generates expanded copies of the unexpanded tfm data structures and keeps them in its memory.

$\text{pdf}\text{T}\text{E}\text{X}$ requires only unexpanded Type 1 font files for font expansion, from which all expanded font versions are internally generated and included (subsetted) into the pdf output file. To enable font expansion, don't forget to set `\pdfadjustspacing` to a value greater than zero.

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 $\text{pdf}\text{T}\text{E}\text{X}$ does not use such huge glyph collections; it uses horizontal scaling instead. This is sub-optimal, and for many fonts, possibly offensive to the design. But, when using pdf, it's not illogical: pdf viewers use so-called Multiple Master fonts when no fonts are embedded and/or can be found on the target system. Such fonts are designed to adapt their design to the different scaling parameters. It is up to the user to determine to what extent mixing slightly remastered fonts can be used without violating the design. Think of an O: when geometrically stretched, the vertical part of the glyph becomes thicker, and looks incompatible with an unscaled original. With a multiple master situation, one can stretch while keeping this thickness compatible.

▶ `\pdfadjustspacing` (integer)

The output that $\text{pdf}\text{T}\text{E}\text{X}$ produces is generally compatible with the normal TEX output: TEX 's typesetting engine is normally unchanged, because the optimization described here is turned off by default. At this moment there are two methods provided.

When `\pdfadjustspacing` is set to 1, stretching is applied *after* T_EX's normal paragraph breaking routines have broken the paragraph into lines. In this case, line breaks are identical to standard T_EX behaviour.

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 T_EX. In fact, paragraphs may even become longer or shorter.

Both alternatives use the extended collection of tfm files that are related to the `stretch` and `shrink` settings as described in the previous section, unless `\pdffontexpand` is given with the option `autoexpand`.

▶ `\efcode` ** (*integer*)

We didn't yet tell the whole story. One can imagine that some glyphs are more sensitive to scaling than others. The `\efcode` primitive can be used to influence the stretchability of a glyph within a given font. The syntax is similar to `\sfcode` (but with the ** required), and defaults to 1000, meaning 100 %.

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

In this example an A may stretch 0.8 times as much as normal and the O is not to be stretched at all. The minimum and maximum stretch is still bound by the font specification, otherwise one would end up with more possible font inclusions than would be comfortable.

▶ `\pdfprotrudechars` (*integer*)

Yet another way of optimizing paragraph breaking is to let certain characters move into the margin. When this primitive is set to 1 (or another positive integer value), the glyphs qualified as such will make this move when applicable. This qualification and the amount of shift are set by the primitives `\rprcode` and `\lprcode`.

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

▶ `\rprcode` ** (*integer*)

The amount that a character from a given font may shift into the right margin ("character protrusion") is set by the primitive `\rprcode`. The protrusion distance is the integer value given to `\rprcode`, multiplied with 0.001 em from the current font. Example:

```
\rptcode\somefont',=200
```

```
\rptcode\somefont'--=150
```

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

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

▶ `\lptcode` ** *<integer>*

This is similar to `\rptcode`, but affects the amount by which characters may protrude into the left margin.

▶ `\pdffontname` ** *<expandable>*

In pdf files produced by pdfT_EX one can recognize a font resource by the prefix `/F` followed by a number, for instance `/F12` or `/F54`. This command returns, for a given T_EX font, the number from the corresponding font resource name. E. g., if `/F12` corresponds to some T_EX font `\foo`, the `\pdffontname\foo` gives the number 12.

In the current implementation, when `\pdfuniquestname` (see below) is set to a positive value, the `\pdffontname` still returns only the number from the font resource name, but not the appended random string.

▶ `\pdffontobjnum` ** *<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 T_EX font `\foo`, the `\pdffontobjnum\foo` gives the number 3.

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

▶ `\pdfincludechars` ** *<general text>*

This command causes pdfT_EX to treat the characters in *<general text>* as if they were used with **, 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.

▶ `\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.

▶ `\pdfmapfile` <map spec>

This primitive is used for managing font map files, to make them known to pdfTeX so that they can be read in. If no `\pdfmapfile` primitive is given, the default map file `pdftex.map` will be read in by pdfTeX.

Normally there is no need for the pdfTeX user to bother about the `\pdfmapfile` primitive, as the main TeX distributions provide nice helper tools that automatically assemble the default font map file. One prominent tool example is the script `updmap` coming with `teTeX`.

The operation mode of the `\pdfmapfile` primitive is selected by a flag letter (+, =, -, or omitted). This flag defines how a map file is going to be handled, and how a collision between an existing map entry and a newer one is resolved; either ignoring a later entry, or replacing or deleting an existing entry. But in any case, map entries of fonts already in use are kept untouched. The companion primitive `\pdfmapline` allows something similar, only that a single map line for one font (instead of a map file name) is given as argument. Here are two examples:

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

The general map handling function is that map items, which are either map file names or single font map lines (in case of the `\pdfmapline` primitive) are put into an auxiliary list of pending map items. During the next page shipout, this list is processed and all pending map items are sequentially scanned for their map entries, and an internal map entry database is updated, using one of the modes described below. Then the list of pending map items is cleared. All `\pdfmapfile` and `\pdfmapline` commands can also be given after shipout of the first page.

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.

`\pdfmapfile{foo.map}` (+/=/- flags omitted) clears the list of pending map items and starts a new list with the only pending file `foo.map`. When the file `foo.map` is scanned, duplicate map entries are ignored and a warning is issued. When

this command is given at the beginning of a T_EX run, the default map file `pdftex.map` will *not* be read in. This is compatible with the former behaviour.

If you want to add support for a new font through an additional font map file while keeping all the existing mappings, don't use this version of command, but instead type either `\pdfmapfile{+myfont.map}` or `\pdfmapfile{=myfont.map}`, as described below.

`\pdfmapfile {+foo.map}` puts the file `foo.map` into the list of pending map items. When the file `foo.map` is scanned, duplicate map entries are ignored and a warning is issued. This is compatible with the former behaviour.

`\pdfmapfile {=foo.map}` puts the file `foo.map` into the list of pending map items. When the file `foo.map` is scanned, matching map entries in the database are replaced by new entries from `foo.map`.

`\pdfmapfile {-foo.map}` puts the file `foo.map` into the list of pending map items. When the file `foo.map` is scanned, matching map entries are deleted from the database.

`\pdfmapfile {}` clears the list of pending map items. It does not affect map entries already registered into the database. This is compatible with the former behaviour. When this command is given at the beginning of a pdfT_EX run, the default map file `pdftex.map` will *not* be read in. This may help with quick program startup, if no fonts are required.

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

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

▶ `\pdfmapline` <map spec>

Similar to `\pdfmapfile`, but here you can give a single map line (like the ones in map files) as an argument. The 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{=...}` variant) allows quick checks of a new font map entry, before writing it into a map file.

`\pdfmapline {}` clears the list of pending map items a similar way as `\pdfmapfile{}` does.

7.4 PDF objects

- ▶ `\pdfobj` [`<object type spec>`]

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, a dictionary object with contents `<general text>` is created.

When however `<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 already before be used by other objects, 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`.

- ▶ `\pdflastobj` (read-only integer)

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

- ▶ `\pdfrefobj` `<object number>`

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.

7.5 XObject forms

The next three primitives support a pdf feature called ‘object reuse’ in pdfTeX. The idea is first to create a XObject form in pdf. The content of this object corresponds to the content of a TeX box; it can contain pictures and references to other XObject form

objects as well. After creation, the XObject form can be used by simply referring to its object number. This feature can be useful for large documents with many similar elements, as it can reduce the duplication of identical objects.

These commands behave similarly to `\pdfobj`, `\pdflastobj` and `\pdfrefobj` but instead of taking raw pdf code, they handle text typeset by TeX.

- ▶ `\pdfxform` [`<attr spec>`] [`<resources spec>`] `<box number>`

This command creates a XObject form corresponding to the contents of the box `<box number>`. The box can contain other raw objects, XObject forms or images as well. It can however *not* contain annotations because they are laid out on a separate layer, are positioned absolutely, and have dedicated housekeeping. `\pdfxform` makes the box void, as `\box` does.

When `<attr spec>` is given, the text will be written as additional attributes of the form. The `<resources spec>` is similar, but the text will be added to the resources dictionary of the form. The text given by `<attr spec>` or `<resources spec>` is written before other keys of the form dictionary and/or the resources dictionary and takes priority over later ones.

- ▶ `\pdfrefxform` `<object number>`

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

- ▶ `\pdflastxform` (read-only integer)

The number of the most recently created XObject form 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 depends on visuals that show up on user request, but are hidden otherwise.

7.6 Graphics inclusion

pdf provides a mechanism for embedding graphic and textual objects: XObject forms. In pdfTeX this mechanism is accessed by means of `\pdfxform`, `\pdflastxform` and `\pdfrefxform`. A special kind of XObjects are bitmap graphics and for manipulating them similar commands are provided.

▶ `\pdfximage` [`<rule spec>`] [`<attr spec>`] [`<page spec>`] `<file spec>`

This command creates an image object. 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 take 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 to the correct resolution. The dimensions of an image can be accessed by enclosing the `\pdfrefximage` command to a box and checking the dimensions of the box:

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

Now we can use `\wd0` and `\ht0` to question the natural size of the image as determined by pdfTeX. When dimensions are specified before the `{somefile.png}`, the graphic is scaled to fit these. Note that, unlike the e. g. `\input` primitive, the filename is supplied between braces.

The image type is specified by the extension of the given file name: `.png` stands for png image, `.jpg` for jpeg, and `.pdf` for pdf file. But once pdfTeX has opened the file, it checks the file type first by looking to 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.

▶ `\pdfrefximage` *(integer)*

The image is kept in memory and will be written to the pdf output only when its number is referred to by `\pdfrefximage` or `\pdfximage` is preceded by `\immediate`. Nothing is appended to the list being built.

`\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 *(integer)* to the pdf output if it has not been written yet.

▶ `\pdflastximage` *(read-only integer)*

The number of the most recently created XObject image is accessible via `\pdflastximage`.

▶ `\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 or jpeg files.

▶ `\pdfimageresolution` *(integer)*

The integer `\pdfimageresolution` parameter (unit: dots per inch, dpi) is a last resort value, used only for bitmap (jpeg, png) 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. Only if no image dimension is 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, finally a default resolution of 72 dpi would be taken. 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 in exif format (unless at least one dimension is stated explicitly); the resolution values coming with exif files are currently ignored.

▶ `\pdfoptionalalwaysusepdfpagebox` *(integer)*

When pdf files are included, the command `\pdfximage` allows selection of which pdf page box to use in the optional field *(image attr spec)*. The integer primitive `\pdfoptionalalwaysusepdfpagebox` allows to globally override this choice by giving them one of the following values: (1) media box, (2) crop box, (3) bleed box, (4) trim box, and (5) artbox.

- ▶ `\pdfoptionpdfinclusionerrorlevel` (integer)

This controls the behaviour of pdfTeX when a pdf file is included that has a newer version than the one specified by this primitive: If it is set to 0, pdfTeX gives only a warning; if it's 1, pdfTeX raises an error.

7.7 Annotations

pdf level 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.

- ▶ `\pdfannot` [`<rule spec>`] `<general text>`

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

- ▶ `\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).

7.8 Destinations and links

The first type of annotation, 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 pdf maintains a dedicated list of these annotations—and some more when optimized—for the sole purpose of speed.

▶ `\pdfdest` \langle dest spec \rangle

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 must be specified in \langle dest type \rangle , which must be one of those mentioned in [table 5](#).

keyword	meaning
<code>fit</code>	fit the page in the window
<code>fitw</code>	fit the width of the page
<code>fitv</code>	fit the height of the page
<code>fitb</code>	fit the ‘Bounding Box’ of the page
<code>fitbh</code>	fit the width of ‘Bounding Box’ of the page
<code>fitbv</code>	fit the height of ‘Bounding Box’ of the page
<code>xyz</code>	goto the current position (see below)

Table 5 Options for display of outline and destinations.

The specification `xyz` can optionally be followed by `zoom` \langle integer \rangle to provide a fixed zoom-in. The \langle integer \rangle is processed like T_EX magnification, i. e. 1000 is the normal page view. When `zoom` \langle integer \rangle is given, the zoom factor changes to number, otherwise the current zoom factor is kept unchanged.

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

▶ `\pdfstartlink` [\langle rule spec \rangle] [\langle attr spec \rangle] \langle action spec \rangle

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. Thus `/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's`.

The `<action spec>` specifies the action that should be performed when the hyperlink is activated while the `<user-action spec>` performs a user-defined action. A typical use of the latter is to specify a url, like `/S /URI /URI (http://www.tug.org/)`, or a named action like `/S /Named /N /NextPage`.

A `<goto-action spec>` performs a GoTo action. Here `<numid>` and `<nameid>` specify the destination identifier (see below). The `<page spec>` specifies the page number of the destination, in this case the zoom factor is given by `<general text>`. A destination can be performed in another pdf file by specifying `<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`. The default behaviour of the `<newwindow spec>` depends on the browser setting.

A `<thread-action spec>` performs an article thread reading. The thread identifier is similar to the destination identifier. A thread can be performed in another pdf file by specifying a `<file spec>`.

▶ `\pdfendlink`

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.

- ▶ `\pdflinkmargin` (dimension)

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

- ▶ `\pdfdestmargin` (dimension)

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

7.9 Bookmarks

- ▶ `\pdfoutline` \langle action spec \rangle [\langle count \langle integer \rangle \rangle] \langle general text \rangle

This primitive creates an outline (or bookmark) entry. The first parameter specifies the action to be taken, and is the same as that allowed for `\pdfstartlink`. The \langle count \rangle specifies the number of direct subentries under this entry; specify 0 or omit it if this entry has no subentries. If the number is negative, then all subentries will be closed and the absolute value of this number specifies the number of subentries. The \langle text \rangle is what will be shown in the outline window. Note that this is limited to characters in the pdf Document Encoding vector. The outline is written to the pdf output immediately.

7.10 Article threads

- ▶ `\pdfthread` \langle rule spec \rangle [\langle attr spec \rangle] \langle id spec \rangle

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

- ▶ `\pdfstartthread` \langle rule spec \rangle [\langle attr spec \rangle] \langle id spec \rangle

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. Note that during output runtime if there are other newly created boxes which have the same nesting level as the vbox/vboxes containing `\pdfstartthread` and `\pdfendthread`, they will be also added into the thread, which is probably

not what you want. To avoid such unconsidered behaviour, it's often enough to wrap boxes that shouldn't belong to the thread by a box to change their box nesting level.

▶ `\pdfendthread`

This ends an article thread started before by `\pdfstartthread`.

▶ `\pdfthreadmargin` (dimension)

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

7.11 Literals and specials

▶ `\pdfliteral` [`direct`] `<general text>`

Like `\special` in normal T_EX, this command inserts raw pdf code into the output. This allows support of color and text transformation. This primitive is heavily used in the METAPOST inclusion macros. Normally pdfT_EX ends a text section in the pdf output and resets the transformation matrix before inserting `<general text>`, however this can be turned off by giving the optional keyword `direct`. This command appends a whatsit node to the list being built. `<general text>` is expanded when the whatsit node is created and not when it is shipped out, as with `\special`.

▶ `\special {pdf: <text> }`

This is equivalent to `\pdfliteral { <text> }`.

▶ `\special {pdf:direct: <text> }`

This is equivalent to `\pdfliteral direct { <text> }`.

7.12 Miscellaneous

▶ `\pdfsavepos`

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 page is finally 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.

▶ `\pdflastxpos` (read-only integer)

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

▶ `\pdflastypos` (read-only integer)

This primitive works similar to `\pdflastxpos`, only it returns the y coordinate.

▶ `\pdftexversion`

Returns the version of pdfTeX multiplied by 100, e. g. for pdfTeX version 1.20b used to produce this document, it returns ‘120’.

▶ `\pdftexrevision`

Returns the revision letter of pdfTeX, e. g. for pdfTeX version 1.20b used to produce this document, it returns letter ‘b’.

▶ `\pdftexbanner`

Returns the pdfTeX banner message, e. g. for the version used here: ‘This is pdfTeX, Version 3.141592-1.20b-2.2 (Web2C 7.5.3) kpathsea version 3.5.3’.

8 Graphics and color

pdfTeX supports inclusion of pictures in png, jpeg, and pdf format. The most common technique with TeX—the inclusion of eps figures—is replaced by pdf inclusion. eps files can be converted to pdf by Ghostscript, Acrobat Distiller or other PostScript-to-pdf converters. The BoundingBox of a pdf file is taken from CropBox if available, otherwise from MediaBox. 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 has been written. The TeXutil utility script and the PStoPDF program that comes with ConTeXt can so a similar job. (Concerning this conversion, they handles complete directories, remove some garbage from files, takes precautions against duplicate conversion, etc.)

Other alternatives for graphics in pdfT_EX are:

L^AT_EX 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 back-end 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, written by Hans Hagen and Tanmoy Bhattacharya) is implemented as a set of macros which reads METAPOST output and supports all of its features.

pdf It is possible to insert arbitrary one-page-only pdf files, with their own fonts and graphics, into a document. The cover page of this document is an example of such an insert, being a one page document generated by pdfT_EX.

For new work, the METAPOST route is highly recommended. For the future, Adobe has announced that they will define a specification for ‘encapsulated pdf’, and this should solve some of the present difficulties.

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

9 Character translation

Characters that are input to pdfT_EX are subject to optional T_EX 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 pdfT_EX. This mapping takes place before the characters enter pdfT_EX’s ‘mouth’. If no tcx file is read, the input characters enter pdfT_EX 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 T_EX character code. If a line contains only one number, characters with this code enter pdfT_EX 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’; remember that character within the range 32..126 are ‘printable’ anyway.

The characters output into the pdf file, e. g. by `\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 `teTeX` manual.

10 Limitations of PDFTeX

pdfTeX currently lacks a colorstack. This can be overcome by the `pdfcolmk` package.

Abbreviations

In this document we used a few abbreviations. For convenience we mention their meaning here.

afm	Adobe Font Metrics
ascii	American Standard Code for Information Interchange
CMacTeX	Macintosh Web2c distribution
ConTeXt	general purpose macro package
dvi	native TeX Device Independent file format
encTeX	encTeX extension to TeX
eps	Encapsulated PostScript
epstopdf	eps to pdf conversion tool

ε -T _E X	an extension to T _E X
exif	Exchangeable Image File format (JPEG file variant)
fpT _E X	Win32 Web2c distribution
hz	Hermann Zapf optimization
jpeg	Joint Photographic Expert Group
L ^A T _E X	general purpose macro package
Mac OS X	Macintosh operating system version 10
METAFont	graphic programming environment, bitmap output
METAPOST	graphic programming environment, vector output
MikT _E X	Win32 distribution
MLT _E X	MLT _E X extension to T _E X
mptopdf	METAPOST to pdf conversion tool
MSDOS	Microsoft DOS platform (Intel)
pdf	Portable Document Format
pdf ε T _E X	ε -T _E X extension producing pdf output
pdfL ^A T _E X	T _E X extension producing pdf output (L ^A T _E X format loaded)
pdfT _E X	T _E X extension producing pdf output
Perl	Perl programming environment
pgc	pdf Glyph Container
pk	Packed bitmap font
png	Portable Network Graphics
PostScript	PostScript
PStoPDF	PostScript to pdf converter (on top of GhostScript)
rgb	Red Green Blue color specification
tcx	T _E X Character Translation
tds	T _E X Directory Standard
teT _E X	T _E X distribution for Unix (based on Web2c)
T _E X	typographic language and program
T _E Xexec	ConT _E Xt command line interface

Texinfo	generate typeset documentation from info pages
T _E X Live	T _E X-Live distribution (multiple platform)
T _E Xutil	ConT _E Xt utility tool
tfm	T _E X Font Metrics
Unix	Unix platform
url	Uniform Resource Locator
web	literate programming environment
Web2c	official multi-platform web environment
Win32	Microsoft Windows platform

Examples of HZ and protruding

In the following sections we will demonstrate pdfT_EX's protruding and hz features, using a text from E. Tufte. This sample text has a lot of punctuation and often needs hyphenation. Former pdfT_EX versions had sometimes problems with combining these features, but from version 1.20b on it should be ok. If you still encounter problems, please try to prepare a small test file that demonstrates the problem and send it to one of the maintainers.

Normal

We thrive in information-thick worlds because of our marvelous and everyday capacity to select, edit, single out, structure, highlight, group, pair, merge, harmonize, synthesize, focus, organize, condense, reduce, boil down, choose, categorize, catalog, classify, list, abstract, scan, look into, idealize, isolate, discriminate, distinguish, screen, pigeonhole, pick over, sort, integrate, blend, inspect, filter, lump, skip, smooth, chunk, average, approximate, cluster, aggregate, outline, summarize, itemize, review, dip into, flip through, browse, glance into, leaf through, skim, refine, enumerate, glean, synopsisize, winnow the wheat from the chaff and separate the sheep from the goats.

We thrive in information-thick worlds because of our marvelous and everyday capacity to select, edit, single out, structure, highlight, group, pair, merge, harmonize, synthesize, focus, organize, condense, reduce, boil down, choose, categorize, catalog, classify, list, abstract, scan, look into, idealize, isolate, discriminate, distinguish, screen, pigeonhole, pick over, sort, integrate, blend, inspect, filter, lump, skip, smooth, chunk, average, approx-

imate, cluster, aggregate, outline, summarize, itemize, review, dip into, flip through, browse, glance into, leaf through, skim, refine, enumerate, glean, synopsise, win-

now the wheat from the chaff and separate the sheep from the goats.

HZ

We thrive in information–thick worlds because of our marvelous and everyday capacity to select, edit, single out, structure, highlight, group, pair, merge, harmonize, synthesize, focus, organize, condense, reduce, boil down, choose, categorize, catalog, classify, list, abstract, scan, look into, idealize, isolate, discriminate, distinguish, screen, pigeonhole, pick over, sort, integrate, blend, inspect, filter, lump, skip, smooth, chunk, average, approximate, cluster, aggregate, outline, summarize, itemize, review, dip into, flip through, browse, glance into, leaf through, skim, refine, enumerate, glean, synopsise, winnow the wheat from the chaff and separate the sheep from the goats.

We thrive in information–thick worlds because of our marvelous and everyday capacity to select, edit, single out, structure, highlight, group, pair, merge, harmonize, synthesize, focus, organize, condense, reduce, boil down, choose, categorize, catalog, classify, list, abstract, scan, look into, idealize, isolate, discriminate, distinguish, screen, pigeonhole,

pick over, sort, integrate, blend, inspect, filter, lump, skip, smooth, chunk, average, approximate, cluster, aggregate, outline, summarize, itemize, review, dip into, flip through, browse, glance into, leaf through, skim, refine, enumerate, glean, synopsise, winnow the wheat from the chaff and separate the sheep from the goats.

Protruding

We thrive in information–thick worlds because of our marvelous and everyday capacity to select, edit, single out, structure, highlight, group, pair, merge, harmonize, synthesize, focus, organize, condense, reduce, boil down, choose, categorize, catalog, classify, list, abstract, scan, look into, idealize, isolate, discriminate, distinguish, screen, pigeonhole, pick over, sort, integrate, blend, inspect, filter, lump, skip, smooth, chunk, average, approximate, cluster, aggregate, outline, summarize, itemize, review, dip into, flip through, browse, glance into, leaf through, skim, refine, enumerate, glean, synopsise, winnow the wheat from the chaff and separate the sheep from the goats.

We thrive in information–thick worlds because of our marvelous and everyday capacity to select, edit, single out, structure, highlight, group, pair, merge, harmonize, synthesize, focus, organize, condense, reduce, boil down, choose, categorize, catalog, classify, list, abstract, scan, look into, idealize, isolate, discriminate, distinguish, screen,

pigeonhole, pick over, sort, integrate, blend, inspect, filter, lump, skip, smooth, chunk, average, approximate, cluster, aggregate, outline, summarize, itemize, review, dip into, flip through, browse, glance into, leaf through, skim, refine, enumerate, glean, synopsisize, winnow the wheat from the chaff and separate the sheep from the goats.

Both

We thrive in information–thick worlds because of our marvelous and everyday capacity to select, edit, single out, structure, highlight, group, pair, merge, harmonize, synthesize, focus, organize, condense, reduce, boil down, choose, categorize, catalog, classify, list, abstract, scan, look into, idealize, isolate, discriminate, distinguish, screen, pigeonhole, pick over, sort, integrate, blend, inspect, filter, lump, skip, smooth, chunk, average, approximate, cluster, aggregate, outline, summarize, itemize, review, dip into, flip through, browse, glance into, leaf through, skim, refine, enumerate, glean, synopsisize, winnow the wheat from the chaff and separate the sheep from the goats.

We thrive in information–thick worlds because of our marvelous and everyday capacity to select, edit, single out, structure, highlight, group, pair, merge, harmonize, synthesize, focus, organize, condense, reduce, boil down, choose, categorize, catalog, classify, list, abstract, scan, look into, idealize, isolate, discriminate, distinguish, screen, pigeonhole,

pick over, sort, integrate, blend, inspect, filter, lump, skip, smooth, chunk, average, approximate, cluster, aggregate, outline, summarize, itemize, review, dip into, flip through, browse, glance into, leaf through, skim, refine, enumerate, glean, synopsisize, winnow the wheat from the chaff and separate the sheep from the goats.

Additional PDF keys

This section is based on the manual on keys written by Martin Schröder, one of the maintainers of pdfTeX.

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. 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 document information dictionary 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. The following key ends up in the document catalog.

key	type	meaning
PTEX.Fullbanner	string	The full version of the binary that produced the file as displayed by <code>pdftex --version</code> , e. g. “This is pdfTeX, Version 3.141592-1.20b-2.2 (Web2C 7.5.4) kpathsea version 3.5.4” This is necessary because the string in the <code>Producer</code> key in the info dictionary is rather short, e. g. “pdfTeX-1.20b”.

Colophon

This manual is typeset in ConTEXt. One can generate an A4 version from the source code by typing:

```
texexec --result=pdfTeX-a.pdf pdfTeX-t
```

Or in letter size:

```
texexec --mode=letter --result=pdfTeX-l.pdf pdfTeX-t
```

Given that the A4 version is typeset, one can generate an A5 booklet by typing:

```
texexec --pdfarrange --paper=a5a4 --print=up --addempty=1,2
--result=pdfptex-b.pdf pdfptex-a
```

Odd and even page sets for non-duplex printers can be generated using `--pages=odd` and `--pages=even` options (which might require some disciplined shuffling of sheet).

This also demonstrates that pdfT_EX can be used for page imposition purposes (given that pdfT_EX and the fonts are set up properly).

GNU Free Documentation License

Version 1.2, November 2002
Copyright © 2000, 2001, 2002
Free Software Foundation, Inc.
59 Temple Place, Suite 330,
Boston, MA 02111-1307 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. Secondly, 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.

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.

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.

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.

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.

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”.

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.

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.

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.

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.

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 <http://www.gnu.org/copyleft/>.

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.