

A Directory Structure for T_EX Files

TUG Working Group on a T_EX Directory Structure (TWG-TDS)

version 1.1 June 23, 2004

Copyright © 1994, 1995, 1996, 1997, 1998, 1999, 2003, 2004 T_EX Users Group.

Permission to use, copy, and distribute this document *without modification* for any purpose and without fee is hereby granted, provided that this notice appears in all copies. It is provided “as is” without expressed or implied warranty.

Permission is granted to copy and distribute modified versions of this document under the conditions for verbatim copying, provided that the modifications are clearly marked and the document is not represented as the official one.

This document is available on any CTAN host (see Appendix D). Please send questions or suggestions by email to tds@tug.org. We welcome all comments. This is version 1.1.

Contents

1	Introduction	2
1.1	History	2
1.2	The role of the TDS	2
1.3	Conventions	3
2	General	3
2.1	Subdirectory searching	3
2.2	Rooting the tree	4
2.3	Local additions	4
2.4	Duplicate filenames	5
3	Top-level directories	5
3.1	Macros	6
3.2	Fonts	8
3.3	Non-font METAFONT files	10
3.4	METAPOST	10
3.5	B _I T _E X	11
3.6	Scripts	11
3.7	Documentation	12
4	Summary	13
4.1	Documentation tree summary	14
A	Unspecified pieces	15
A.1	Portable filenames	15
B	Implementation issues	16
B.1	Adoption of the TDS	16
B.2	More on subdirectory searching	17
B.3	Example implementation-specific trees	17
C	Is there a better way?	19
C.1	Macro structure	19
C.2	Font structure	20
C.3	Documentation structure	21
D	Related references	21
E	Contributors	22

1 Introduction

\TeX is a powerful, flexible typesetting system used by many people around the world. It is extremely portable and runs on virtually all operating systems. One unfortunate side effect of \TeX 's flexibility, however, is that there has been no single "right" way to install it. This has resulted in many sites having different installed arrangements.

The primary purpose of this document is to describe a standard \TeX Directory Structure (TDS): a directory hierarchy for macros, fonts, and the other implementation-independent \TeX system files. As a matter of practicality, this document also suggests ways to incorporate the rest of the \TeX files into a single structure. The TDS has been designed to work on all modern systems. In particular, the Technical Working Group (TWG) believes it is usable under MacOS, MS-DOS, OS/2, Unix, VMS, and Windows NT. We hope that administrators and developers of both free and commercial \TeX implementations will adopt this standard.

This document is intended both for the \TeX system administrator at a site and for people preparing \TeX distributions—everything from a complete runnable system to a single macro or style file. It may also help \TeX users find their way around systems organized this way. It is not a tutorial: we necessarily assume knowledge of the many parts of a working \TeX system. If you are unfamiliar with any of the programs or file formats we refer to, consult the references in Appendix D.

1.1 History

Version 1.0 of the TDS was released in February 2003.

Version 1.1 was released in June 2004, with the following non-editorial changes:

- Inputs for \TeX extensions included under `tex`, instead of in their own top-level directories (Section 3.1.1)
- New top-level directory `scripts` (Section 3.6).
- New subdirectories `lig`, `opentype`, `truetype`, and `type3` under `fonts` (Section 3.2).
- `enc`, `lig`, and `map` all use $\langle syntax \rangle / \langle package \rangle$ subdirectories (Section 3.2).
- `pfm` files specified to go under `type1`, and `inf` files under `afm` (Section 3.2).

1.2 The role of the TDS

The role of the TDS is to stabilize the organization of \TeX -related software packages that are installed and in use, possibly on multiple platforms simultaneously.

At first glance, it may seem that the Comprehensive \TeX Archive Network (CTAN) fulfills at least part of this role, but this is not the case. The role of CTAN is to simplify archiving and distribution, not installation and use.

In fact, the roles of the TDS and CTAN are frequently in conflict, as we will see. For distribution, many different types of files must be combined into a single unit; for use, it is traditional to segregate files (even similar files) from a single package into separate, occasionally distant, directories.

1.3 Conventions

In this document, “/” is used to separate filename components; for example, `texmf/fonts`. This is the Unix convention but the ideas are in no way Unix-specific.

In this document, “ \TeX ” generally means the \TeX system, including METAFONT, DVI drivers, utilities, etc., not just the \TeX program itself.

The word “package” in this document has its usual meaning: a set of related files distributed, installed, and maintained as a unit. This is *not* a $\LaTeX 2_{\epsilon}$ package, which is a style file supplementing a document class.

We use the following typographic conventions:

literal Literal text such as `filename` is typeset in typewriter type.

<replaceable> Replaceable text such as *<package>*, identifying a class of things, is typeset in italics inside angle brackets.

2 General

This section describes common properties throughout the TDS tree.

2.1 Subdirectory searching

Older \TeX installations store large numbers of related files in single directories, for example, all TFM files and/or all \TeX input files.

This monolithic arrangement hinders maintenance of a \TeX system: it is difficult to determine what files are used by what packages, what files need to be updated when a new version is installed, or what files should be deleted if a package is removed. It is also a source of error if two or more packages happen to have input files with the same name.

Therefore, the TWG felt each package should be in a separate directory. But we recognized that explicitly listing all directories to be searched would be unbearable. A site may wish to install dozens of packages. Aside from anything else, listing that many directories would produce search paths many thousands of characters long, overflowing the available space on some systems.

Also, if all directories are explicitly listed, installing or removing a new package would mean changing a path as well as installing or removing the actual files. This would be a time-consuming and error-prone operation, even with implementations that provide some way to specify the directories to search at runtime. On systems without runtime configuration, it would require recompiling software, an intolerable burden.

As a result, the TWG concluded that a comprehensive TDS requires implementations to support some form of implicit subdirectory searching. More precisely, implementations must make it possible to specify that \TeX , METAFONT, and their companion utilities search in both a specified directory and recursively through all subdirectories of that directory when looking for an input file. Other forms of subdirectory searching, for example recursive-to-one-level searches, may also be provided. We encourage implementors to provide subdirectory searching at the option of the installer and user for all paths.

The TDS does not specify a syntax for specifying recursive searching, but we encourage implementors to provide interoperability (see Section B.2).

2.2 Rooting the tree

In this document, we shall designate the root TDS directory by ‘`texmf`’ (for “`TEX` and `META-FONT`”). We recommend using that name where possible, but the actual name of the directory is up to the installer. On PC networks, for example, this could map to a logical drive specification such as `T:`.

Similarly, the location of this directory on the system is site-dependent. It may be at the root of the file system; on Unix systems, `/usr/local/share`, `/usr/local`, `/usr/local/lib`, and `/opt` are common choices.

The name `texmf` was chosen for several reasons: it reflects the fact that the directory contains files pertaining to an entire `TEX` system (including `METAFONT`, `METAPOST`, `BIBTEX`, etc.), not just `TEX` itself; and it is descriptive of a generic installation rather than a particular implementation.

A site may choose to have more than one TDS hierarchy installed (for example, when installing an upgrade). This is perfectly legitimate.

2.3 Local additions

The TDS cannot specify precisely when a package is or is not a “local addition”. Each site must determine this according to its own conventions. At the two extremes, one site might wish to consider “nonlocal” all files not acquired as part of the installed `TEX` distribution; another site might consider “local” only those files that were actually developed at the local site and not distributed elsewhere.

We recognize two common methods for local additions to a distributed `texmf` tree. Both have their place; in fact, some sites employ both simultaneously:

1. A completely separate tree which is a TDS structure itself; for example, `/usr/local/umbtex` at the University of Massachusetts at Boston. This is another example of the multiple `texmf` hierarchies mentioned in the previous section.
2. A directory named ‘`local`’ at any appropriate level, for example, in the `<format>`, `<package>`, and `<supplier>` directories discussed in the following sections. The TDS reserves the directory name `local` for this purpose.

We recommend using `local` for site-adapted configuration files, such as `language.dat` for the Babel package or `graphics.cfg` for the graphics package. Unmodified configuration files from a package should remain in the package directory. The intent is to separate locally modified or created files from distribution files, to ease installing new releases.

One common case of local additions is dynamically generated files, e.g., PK fonts by the `mktexpk` script (which originated in `Dvips` as `MakeTXPk`). A site may store the generated files directly in any of:

- their standard location in the main TDS tree (if it can be made globally writable);
- an alternative location in the main TDS tree (for example, under `texmf/fonts/tmp`);
- a second complete TDS tree (as outlined above);
- any other convenient directory (perhaps under `/var`, for example `/var/spool/fonts`).

No one solution will be appropriate for all sites.

2.4 Duplicate filenames

Different files by the same name may exist in a TDS tree. The TDS generally leaves unspecified which of two files by the same name in a search path will be found, so generally the only way to reliably find a given file is for it to have a unique name. However, the TDS requires implementations to support the following exceptions:

- Names of $\text{T}_{\text{E}}\text{X}$ input files must be unique within each first-level subdirectory of `texmf/tex` and `texmf/tex/generic`, but not within all of `texmf/tex`; i.e., different $\text{T}_{\text{E}}\text{X}$ formats may have files by the same name. (Section 3.1 discusses this further.) Thus, no single format-independent path specification, such as a recursive search beginning at `texmf/tex` specifying no other directories, suffices. So implementations must provide format-dependent path specifications, for example via wrapper scripts or configuration files.
- Many font files will have the same name (e.g., `cmr10.pk`), as discussed in Section 3.2.2. Implementations must distinguish these files by mode and resolution.

All implementations we know of already have these capabilities.

One place where duplicate names are likely to occur is not an exception:

- Names of METAFONT input files (as opposed to bitmaps) must be unique within all of `texmf/fonts`. In practice, this is a problem with some variants of Computer Modern which contain slightly modified files named `punct.mf`, `roman1.mf`, and so on. We believe the only feasible solution is to rename the derivative files to be unique.

3 Top-level directories

The directories under the `texmf` root identify the major components of a $\text{T}_{\text{E}}\text{X}$ system (see Section 4 for a summary). A site may omit any unneeded directories.

Although the TDS by its nature can specify precise locations only for implementation-independent files, we recognize that installers may well wish to place other files under `texmf` to simplify administration of the $\text{T}_{\text{E}}\text{X}$ tree, especially if it is maintained by someone other than the system administrator. Therefore, additional top-level directories may be present.

The top-level directories specified by the TDS are:

`tex` for $\text{T}_{\text{E}}\text{X}$ files (Section 3.1).

`fonts` for font-related files (Section 3.2).

`metafont` for METAFONT files which are not fonts (Section 3.3).

`metapost` for METAPOST files (Section 3.4).

`bibtex` for $\text{BIB}_{\text{E}}\text{X}$ files (Section 3.5).

`scripts` for platform-independent executables (Section 3.6).

`doc` for user documentation (Section 3.7).

`source` for sources. This includes both traditional program sources (for example, Web2C sources go in `texmf/source/web2c`) and, e.g., $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$ `dtx` sources (which go in `texmf/source/latex`). The TDS leaves unspecified any structure under `source`.

`source` is intended for files which are not needed at runtime by any \TeX program; it should not be included in any search path. For example, `plain.tex` does not belong under `texmf/source`, even though it is a “source file” in the sense of not being derived from another file. (It goes in `texmf/tex/plain/base`, as explained in Section 3.1).

⟨implementation⟩ for implementations (examples: `emtex`, `vtex`, `web2c`), to be used for whatever purpose deemed suitable by the implementor or \TeX administrator. That is, files that cannot be shared between implementations, such as pool files (`tex.pool`) and memory dump files (`plain.fmt`) go here, in addition to implementation-wide configuration files. See Section B.3 for examples of real *⟨implementation⟩* trees.

Such implementation-specific configuration files should *not* be located using the main \TeX input search path (e.g., `TEXINPUTS`). This must be reserved for files actually read by a \TeX engine. See Section 3.1.1.

⟨program⟩ for program-specific input and configuration files for any \TeX -related programs (examples: `mft`, `dvips`). In fact, the `tex`, `metafont`, `metapost`, and `bibtex` items above may all be seen as instances of this case.

3.1 Macros

\TeX macro files shall be stored in separate directories, segregated by \TeX format and package name (we use ‘format’ in its traditional \TeX sense to mean a usefully `\dump`-able package):

`texmf/tex/⟨format⟩/⟨package⟩/`

⟨format⟩ is a format name (examples: `amstex`, `latex`, `plain`, `texinfo`).

The TDS allows distributions that can be used as either formats or packages (e.g., `Texinfo`, `Eplain`) to be stored at either level, at the option of the format author or \TeX administrator. We recommend that packages used as formats at a particular site be stored at the *⟨format⟩* level: by adjusting the \TeX inputs search path, it will be straightforward to use them as macro packages under another format, whereas placing them in another tree completely obscures their use as a format.

The TDS reserves the following *⟨format⟩* names:

- `generic`, for input files that are useful across a wide range of formats (examples: `null.tex`, `path.sty`). Generally, this means any format that uses the category codes of Plain \TeX and does not rely on any particular format. This is in contrast to those files which are useful only with Plain \TeX (which go under `texmf/tex/plain`), e.g., `testfont.tex` and `plain.tex` itself.
- `local`, for local additions. See Section 2.3.

Thus, for almost every format, it is necessary to search at least the *⟨format⟩* directory and then the `generic` directory (in that order). Other directories may need to be searched as well, depending on the format. When using $\mathcal{A}\mathcal{M}\mathcal{S}$ - \TeX , for example, the `amstex`, `plain`, and `generic` directories should be searched, because $\mathcal{A}\mathcal{M}\mathcal{S}$ - \TeX is compatible with Plain.

⟨package⟩ is a \TeX package name (examples: `babel`, `texdraw`).

In the case where a format consists of only a single file and has no auxiliary packages, that file can simply be placed in the *⟨format⟩* directory, instead of *⟨format⟩/base*. For

example, Texinfo may go in `texmf/tex/texinfo/texinfo.tex`, not `texmf/tex/texinfo/base/texinfo.tex`.

The TDS reserves the following *⟨package⟩* names:

- **base**, for the base distribution of each format, including files used by INITEX when dumping format files. For example, in the standard L^AT_EX distribution, the `ltx` files created during the build process. Another example: the `.ini` driver files for formats used by T_EX Live and other distributions.
- **hyphen**, for hyphenation patterns, including the original American English `hyphen.tex`. These are typically used only by INITEX. In most situations, this directory need exist only under the **generic** format.
- **images**, for image input files, such as Encapsulated PostScript figures. Although it is somewhat non-intuitive for these to be under a directory named “**tex**”, T_EX needs to read these files to glean bounding box or other information. A mechanism for sharing image inputs between T_EX and other typesetting programs (e.g., Interleaf, FrameMaker) is beyond the scope of the TDS. In most situations, this directory need exist only under the **generic** format.
- **local**, for local additions and configuration files. See Section 2.3.
- **misc**, for packages that consist of a single file. An administrator or package maintainer may create directories for single-file packages at their discretion, instead of using **misc**.

3.1.1 Extensions

T_EX has spawned many companion and successor programs (“engines”), such as PDFT_EX, Omega, and others. The TDS specifies that the input files for such programs (using a T_EX-like syntax) be placed within the top-level **tex** directory, either at the top level or within a format subdirectory, even though the original T_EX program may not be able to read them. For example:

```
texmf/tex/aleph
texmf/tex/enctex
```

This is a change from TDS 1.0, which specified top-level *⟨extension⟩* directories for each such program. We felt the new approach is preferable, because:

- Authors of relevant packages typically make their code detect the engine being used, and issue error messages or adapt to circumstances appropriately. Furthermore, as a package matures, it may support multiple engines. Thus, a package could conceivably be placed in any of several top-level directories, at different times. Putting all packages under the top-level **tex** directory provides a stable location over time.
- Users need to be able to switch between engines, and configuring different search paths for each engine is difficult and error-prone.

Thus, in practice, having different top-level directories caused difficulties for everyone involved—users, package authors, site administrators, and system distributors.

Please contrast this approach with the *⟨implementation⟩* top-level subdirectory (Section 3), which is to be used for configuration files that (presumably) do not use T_EX syntax and in any case should not be found along the main T_EX input search path.

3.2 Fonts

Font files are stored in separate directories, segregated by file type, and then (in most cases) font supplier and typeface. PK and GF files need additional structure, as detailed in the next section.

```
texmf/fonts/⟨type⟩/⟨supplier⟩/⟨typeface⟩/  
texmf/fonts/enc,lig,map/⟨subpath⟩/
```

⟨*type*⟩ is the type of font file. The TDS reserves the following ⟨*type*⟩ names for common T_EX file types:

- **afm**, for Adobe font metrics, and **inf** files.
- **gf**, for generic font bitmap files.
- **opentype**, for OpenType fonts.
- **pk**, for packed bitmap files.
- **source**, for font sources (METAFONT files, property lists, etc.).
- **tfm**, for T_EX font metric files.
- **truetype**, for TrueType fonts.
- **type1**, for PostScript Type 1 fonts (in **pfa**, **pfb**, or any other format), and **pfm** files.
- **type3**, for PostScript Type 3 fonts.
- **vf**, for virtual fonts.

The TDS also reserves the names **enc**, **lig**, and **map** for font encoding, ligature, and mapping files, respectively. All of these directories are structured the same way, with ⟨*syntax*⟩ subdirectories, and then ⟨*package*⟩ subsubdirectories. Each of these file types is intended to be searched along its own recursively-searched path. The names of the actual files must be unique within their subtree, as usual. Examples:

```
fonts/map/dvipdfm/updmap/dvipdfm.map  
fonts/map/dvips/lm/lm.map  
fonts/enc/dvips/base/8r.enc
```

The Fontname and Dvips packages have more examples of the **enc** and **map** types. The **afm2pl** program uses **lig** files.

pfm files are included in the **type1** directory, instead of being given their own directory, for two reasons: 1) a **.pfm** file is always an adjunct to a given **.pfb** file; 2) they must be installed from the same directory for Windows programs other than T_EX to use them.

inf files are included in the **afm** directory, since an **inf** and **afm** file can be used to generate a **pfm**. (Unfortunately, Adobe Type Manager and perhaps other software requires that the **pfb** be in the same directory as **afm** and **inf** for installation.)

As usual, a site may omit any of these directories that are unnecessary. **gf** is a particularly likely candidate for omission.

⟨*supplier*⟩ is a name identifying font source (examples: **adobe**, **ams**, **public**). The TDS reserves the following ⟨*supplier*⟩ names:

- **ams**, for the American Mathematical Society's $\mathcal{A}\mathcal{M}\mathcal{S}$ -fonts collection.
- **local**, for local additions. See Section 2.3.
- **public**, for freely redistributable fonts where the supplier neither (1) requested their own directory (e.g., **ams**), nor (2) also made proprietary fonts (e.g., **adobe**). It does

not contain all extant freely distributable fonts, nor are all files therein necessarily strictly public domain.

- `tmp`, for dynamically-generated fonts, as is traditional on some systems. It may be omitted if unnecessary, as usual.

`<typeface>` is the name of a typeface family (examples: `cm`, `euler`, `times`). The TDS reserves the following `<typeface>` names:

- `cm` (within `public`), for the 75 fonts defined in *Computers and Typesetting, Volume E*.
- `latex` (within `public`), for those fonts distributed with L^AT_EX in the base distribution.
- `local`, for local additions. See Section 2.3.

Some concrete examples:

```
texmf/fonts/source/public/pandora/pnr10.mf
texmf/fonts/tfm/public/cm/cmr10.tfm
texmf/fonts/type1/adobe/utopia/putr.pfa
```

For complete supplier and typeface name lists, consult *Filenames for T_EX fonts* (see Appendix D).

3.2.1 Font bitmaps

Font bitmap files require two characteristics in addition to the above to be uniquely identifiable: (1) the type of device (i.e., mode) for which the font was created; (2) the resolution of the bitmap.

Following common practice, the TDS segregates fonts with different device types into separate directories. See `modes.mf` in Appendix D for recommended mode names.

Some printers operate at more than one resolution (e.g., at 300 dpi and 600 dpi), but each such resolution will necessarily have a different mode name. Nothing further is needed, since implicit in the T_EX system is the assumption of a single target resolution.

Two naming strategies are commonly used to identify the resolution of bitmap font files. On systems that allow long filenames (and in the original METAFONT program itself), the resolution is included in the filename (e.g., `cmr10.300pk`). On systems which do not support long filenames, fonts are generally segregated into directories by resolution (e.g., `dpi300/cmr10.pk`).

Because the TDS cannot require long filenames, we must use the latter scheme for naming fonts. So we have two more subdirectory levels under `pk` and `gf`:

```
texmf/fonts/pk/<mode>/<supplier>/<typeface>/dpi<nnn>/
texmf/fonts/gf/<mode>/<supplier>/<typeface>/dpi<nnn>/
```

`<mode>` is a name which identifies the device type (examples: `cx`, `ljfour`, `modeless`). Usually, this is the name of the METAFONT mode used to build the PK file. For fonts rendered as bitmaps by a program that does not distinguish between different output devices, the `<mode>` name shall be simply `modeless`. The `<mode>` level shall not be omitted, even if only a single mode happens to be in use.

`dpi<nnn>` specifies the resolution of the font (examples: `dpi300`, `dpi329`). ‘dpi’ stands for dots per inch, i.e., pixels per inch. We recognize that pixels per millimeter is used in many parts of the world, but dpi is too traditional in the T_EX world to consider changing now.

The integer $\langle nnn \rangle$ is to be calculated as if using METAFONT arithmetic and then rounded; i.e., it is the integer METAFONT uses in its output `gf` filename. We recognize small differences in the resolution are a common cause of frustration among users, however, and recommend implementors follow the level 0 DVI driver standard (see Appendix D) in bitmap font searches by allowing a fuzz of $\pm 0.2\%$ (with a minimum of 1) in the $\langle dpi \rangle$.

Implementations may provide extensions to the basic naming scheme, such as long filenames (as in the original METAFONT) and font library files (as in `emTeX`'s `.fli` files), provided that the basic scheme is also supported.

3.2.2 Valid font bitmaps

The TWG recognizes that the use of short filenames has many disadvantages. The most vexing is that it results in the creation of dozens of different files with the same name. At a typical site, `cmr10.pk` will be the filename for Computer Modern Roman 10 pt at 5–10 magnifications for 2–3 modes. (Section 2.4 discusses duplicate filenames in general.)

To minimize this problem, we strongly recommend that PK files contain enough information to identify precisely how they were created: at least the mode, base resolution, and magnification used to create the font.

This information is easy to supply: a simple addition to the local modes used for building the fonts with METAFONT will automatically provide the required information. If you have been using a local modes file derived from (or that is simply) `modes.mf` (see Appendix D), the required information is already in your PK files. If not, a simple addition based on the code found in `modes.mf` can be made to your local modes file and the PK files rebuilt.

3.3 Non-font METAFONT files

Most METAFONT input files are font programs or parts of font programs and are thus covered by the previous section. However, a few non-font input files do exist. Such files shall be stored in:

`texmf/metafont/ $\langle package \rangle$ /`

$\langle package \rangle$ is the name of a METAFONT package (for example, `mfpic`).

The TDS reserves the following $\langle package \rangle$ names:

- `base`, for the standard METAFONT macro files as described in *The METAFONTbook*, such as `plain.mf` and `expr.mf`.
- `local`, for local additions. See Section 2.3.
- `misc`, for METAFONT packages consisting of only a single file (for example, `modes.mf`). An administrator or package maintainer may create directories for single-file packages at their discretion, instead of using `misc`.

3.4 METAPOST

METAPOST is a picture-drawing language developed by John Hobby, derived from Knuth's METAFONT. Its primary purpose is to output Encapsulated POSTSCRIPT instead of bitmaps.

METAPOST input files and the support files for METAPOST-related utilities shall be stored in:

`texmf/metapost/⟨package⟩/`

`⟨package⟩` is the name of a METAPOST package. At the present writing none exist, but the TWG thought it prudent to leave room for contributed packages that might be written in the future.

The TDS reserves the following `⟨package⟩` names:

- **base**, for the standard METAPOST macro files, such as `plain.mp`, `mfplain.mp`, `boxes.mp`, and `graph.mp`. This includes files used by INIMP when dumping mem files containing preloaded macro definitions.
- **local**, for local additions. See Section 2.3.
- **misc**, for METAPOST packages consisting of only a single file. An administrator or package maintainer may create directories for single-file packages at their discretion, instead of using **misc**.
- **support**, for additional input files required by METAPOST utility programs, including a font map, a character adjustment table, and a subdirectory containing low-level METAPOST programs for rendering some special characters.

3.5 B_IB_TE_X

B_IB_TE_X-related files shall be stored in:

`texmf/bibtex/bib/⟨package⟩/`
`texmf/bibtex/bst/⟨package⟩/`

The **bib** directory is for B_IB_TE_X database (`.bib`) files, the **bst** directory for style (`.bst`) files.

`⟨package⟩` is the name of a B_IB_TE_X package. The TDS reserves the following `⟨package⟩` names (the same names are reserved under both **bib** and **bst**):

- **base**, for the standard B_IB_TE_X databases and styles, such as `xampl.bib`, `plain.bst`.
- **local**, for local additions. See Section 2.3.
- **misc**, for B_IB_TE_X packages consisting of only a single file. An administrator or package maintainer may create directories for single-file packages at their discretion, instead of using **misc**.

3.6 Scripts

The top-level **scripts** directory is for platform-independent executables, such as Perl, Python, and shell scripts, and Java class files. Subdirectories under **scripts** are package names. This eases creating distributions, by providing a common place for such platform-independent programs.

The intent is not for all such directories to be added to a user's command search path, which would be quite impractical. Rather, these executables are primarily for the benefit of wrapper scripts in whatever executable directory a distribution may provide (which is not specified by the TDS).

Truly auxiliary scripts which are invoked directly by other programs, rather than wrapper scripts, may also be placed here. That is, **scripts** also serves as a platform-independent analog of the standard Unix `libexec` directory.

We recommend using extensions specifying the language (such as `.pl`, `.py`, `.sh`) on these files, to help uniquely identify the name. Since the intent of the TDS is for programs in `scripts` not to be invoked directly by users, this poses no inconvenience.

For example, in the T_EX Live distribution, the ConT_EXt user-level program `texexec` can exist as a small wrapper script in each `bin/⟨platform⟩/texexec` (which is outside the `texmf` tree), which merely finds and calls `texmf/scripts/context/perl/texexec.pl`.

Examples:

```
scripts/context/perl/texexec.pl
scripts/context/ruby/example.rb
scripts/thumbpdf/thumbpdf.pl
```

The TDS does not specify a location for platform-dependent binary executables, whether auxiliary or user-level.

3.7 Documentation

Most packages come with some form of documentation: user manuals, example files, programming guides, etc. In addition, many independent files not part of any macro or other package have been created to describe various aspects of the T_EX system.

The TDS specifies that these additional documentation files shall be stored in a structure that parallels to some extent the `fonts` and `tex` directories, as follows:

```
texmf/doc/⟨category⟩/...
```

⟨*category*⟩ identifies the general topic of documentation that resides below it; for example, a T_EX format name (`latex`), program name (`bibtex`, `tex`), language (`french`, `german`), a file format (`info`, `man`), or other system components (`web`, `fonts`).

One possible arrangement is to organize `doc` by language, with all the other category types below that. This helps users find documentation in the language(s) in which they are fluent. Neither this nor any other particular arrangement is required, however.

Within each ⟨*category*⟩ tree for a T_EX format, the directory `base` is reserved for base documentation distributed by the format's maintainers.

The TDS reserves the following category names:

- `general`, for standalone documents not specific to any particular program (for example, Joachim Schrod's *Components of T_EX*).
- `help`, for meta-information, such as FAQ's, the T_EX Catalogue, etc.
- `info`, for processed Texinfo documents. (Info files, like anything else, may also be stored outside the TDS, at the installer's option.)
- `local`, for local additions. See Section 2.3.

The `doc` directory is intended for implementation-independent and operating system-independent documentation files. Implementation-dependent files are best stored elsewhere, as provided for by the implementation and/or T_EX administrator (for example, VMS help files under `texmf/vms/help`).

The documentation directories may contain T_EX sources, DVI files, POSTSCRIPT files, text files, example input files, or any other useful documentation format(s).

See Section 4.1 for a summary.

4 Summary

A skeleton of a TDS `texmf` directory tree. This is not to imply these are the only entries allowed. For example, `local` may occur at any level.

<code>bibtex/</code>	BiBTeX input files
<code>bib/</code>	BiBTeX databases
<code>base/</code>	base distribution (e.g., <code>xampl.bib</code>)
<code>misc/</code>	single-file databases
<code><package>/</code>	name of a package
<code>bst/</code>	BiBTeX style files
<code>base/</code>	base distribution (e.g., <code>plain.bst</code> , <code>acm.bst</code>)
<code>misc/</code>	single-file styles
<code><package>/</code>	name of a package
<code>doc/</code>	see Section 3.7 and the summary below
<code>fonts/</code>	font-related files
<code><type>/</code>	file type (e.g., <code>pk</code>)
<code><mode>/</code>	type of output device (for <code>pk</code> and <code>gf</code> only)
<code><supplier>/</code>	name of a font supplier (e.g., <code>public</code>)
<code><typeface>/</code>	name of a typeface (e.g., <code>cm</code>)
<code>dpi<nnn>/</code>	font resolution (for <code>pk</code> and <code>gf</code> only)
<code><implementation>/</code>	TeX implementations, by name (e.g., <code>emtex</code>)
<code>local/</code>	files created or modified at the local site
<code>metafont/</code>	METAFONT (non-font) input files
<code>base/</code>	base distribution (e.g., <code>plain.mf</code>)
<code>misc/</code>	single-file packages (e.g., <code>modes.mf</code>)
<code><package>/</code>	name of a package (e.g., <code>mfpic</code>)
<code>metapost/</code>	METAPOST input and support files
<code>base/</code>	base distribution (e.g., <code>plain.mp</code>)
<code>misc/</code>	single-file packages
<code><package>/</code>	name of a package
<code>support/</code>	support files for METAPOST-related utilities
<code>mft/</code>	MFT inputs (e.g., <code>plain.mft</code>)
<code><program>/</code>	TeX-related programs, by name (e.g., <code>dvips</code>)
<code>source/</code>	program source code by name (e.g., <code>latex</code> , <code>web2c</code>)
<code>tex/</code>	TeX input files
<code><engine>/</code>	name of an engine (e.g., <code>aleph</code>); can also be lower
<code><format>/</code>	name of a format (e.g., <code>plain</code>)
<code>base/</code>	base distribution for format (e.g., <code>plain.tex</code>)
<code>misc/</code>	single-file packages (e.g., <code>webmac.tex</code>)
<code>local/</code>	local additions to or local configuration files for <code><format></code>
<code><package>/</code>	name of a package (e.g., <code>graphics</code> , <code>mfnfss</code>)
<code>generic/</code>	format-independent packages
<code>hyphen/</code>	hyphenation patterns (e.g., <code>hyphen.tex</code>)
<code>images/</code>	image input files (e.g., Encapsulated PostScript)
<code>misc/</code>	single-file format-independent packages (e.g., <code>null.tex</code>).
<code><package>/</code>	name of a package (e.g., <code>babel</code>)

4.1 Documentation tree summary

An example skeleton of a TDS directory tree under `texmf/doc`. This is not to imply these are the only entries allowed, or that this structure must be followed precisely for the entries listed.

As mentioned, the `texmf/doc` tree may be organized by language, so that all documentation in French, say, is in a `french` subdirectory. In that case, the example structure here would be in a given language directory.

```
ams/
  amsfonts/  amsfonts.faq, amfndoc
  amslatex/  amslatex.faq, amsldoc
  amstex/    amsguide, joyerr
bibtex/     BibTEX
  base/      btxdoc.tex
fonts/
  fontname/  Filenames for TEX fonts
  oldgerm/   corkpapr
⟨format⟩/   name of a TEX format (e.g., generic, latex)
  base/      for the base distribution
  misc/      for contributed single-file package documentation
  ⟨package⟩/ for package
general/    across programs, generalities
  errata/    errata, errata[1-8]
  texcomp/   Components of TEX
help/       meta-information
  ctan/      info about CTAN mirror sites
  faq/       FAQs of comp.text.tex, etc.
info/       GNU Info files, made from Texinfo sources
latex/      example of ⟨format⟩
  base/      ltnews*, *guide, etc.
  graphics/  grfguide
local/      site-specific documentation
man/        Unix man pages
⟨program⟩/  TEX-related programs, by name (examples follow)
metafont/   mfbook.tex, metafont-for-beginners, etc.
metapost/   mpman, manfig, etc.
tex/        texbook.tex, A Gentle Introduction to TEX, etc.
web/        webman, cwebman
```

A Unspecified pieces

The TDS cannot address the following aspects of a functioning \TeX system:

1. The location of executable programs: this is too site-dependent even to recommend a location, let alone require one. A site may place executables outside the `texmf` tree altogether (e.g., `/usr/local/bin`), in a platform-dependent directory within `texmf`, or elsewhere.
2. Upgrading packages when new releases are made: we could find no way of introducing version specifiers into `texmf` that would do more good than harm, or that would be practical for even a plurality of installations.
3. The location of implementation-specific files (e.g., \TeX `.fmt` files): by their nature, these must be left to the implementor or \TeX maintainer. See Section B.3.
4. Precisely when a package or file should be considered “local”, and where such local files are installed. See Section 2.3 for more discussion.

A.1 Portable filenames

The TDS cannot require any particular restriction on filenames in the tree, since the names of many existing \TeX files conform to no standard scheme. For the benefit of people who wish to make a portable \TeX distribution or installation, however, we outline here the necessary restrictions. The TDS specifications themselves are compatible with these.

ISO-9660 is the only universally acceptable file system format for CD-ROMs. A subset thereof meets the stringent limitations of all operating systems in use today. It specifies the following:

- File and directory names, not including any directory path or extension part, may not exceed eight characters.
- Filenames may have a single extension. Extensions may not exceed three characters. Directory names may not have an extension.
- Names and extensions may consist of *only* the characters A–Z, 0–9, and underscore. Lowercase letters are excluded.
- A period separates the filename from the extension and is always present, even if the name or extension is missing (e.g., `FILENAME.` or `.EXT`).
- A version number, ranging from 1–32767, is appended to the file extension, separated by a semicolon (e.g., `FILENAME.EXT;1`).
- Only eight directory levels are allowed, including the top-level (mounted) directory (see Section 2.2). Thus, the deepest valid ISO-9660 path is:

```
texmf/L2/L3/L4/L5/L6/L7/L8/FOO.BAR;1
1      2 3 4 5 6 7 8
```

The deepest TDS path needs only seven levels:

```
texmf/fonts/pk/cx/public/cm/dpi300/cmr10.pk
1      2      3 4 5      6 7
```

Some systems display a modified format of ISO-9660 names, mapping alphabetic characters to lowercase, removing version numbers and trailing periods, etc.

Before the December 1996 release, \LaTeX used mixed-case names for font descriptor files. Fortunately, it never relied on case alone to distinguish among the files. Nowadays, it uses only monospace names.

B Implementation issues

We believe that the TDS can bring a great deal of order to the current anarchic state of many \TeX installations. In addition, by providing a common frame of reference, it will ease the burden of documenting administrative tasks. Finally, it is a necessary part of any reasonable system of true “drop-in” distribution packages for \TeX .

B.1 Adoption of the TDS

[This section is retained for historical purposes; the TDS is now quite firmly entrenched in most \TeX distributions.]

We recognize that adoption of the TDS will not be immediate or universal. Most \TeX administrators will not be inclined to make the final switch until:

- Clear and demonstrable benefits can be shown for the TDS.
- TDS-compliant versions of all key programs are available in ported, well-tested forms.
- A “settling” period has taken place, to flush out problems. The public release of the first draft of this document was the first step in this process.

Consequently, most of the first trials of the TDS will be made by members of the TDS committee and/or developers of \TeX -related software. This has already taken place during the course of our deliberations (see Appendix D for a sample tree available electronically). They will certainly result in the production of a substantial number of TDS-compliant packages. Indeed, the $\text{te}\text{\TeX}$ and \TeX Live distributions are TDS-compliant and in use now at many sites.

Once installable forms of key TDS-compliant packages are more widespread, some \TeX administrators will set up TDS-compliant trees, possibly in parallel to existing production directories. This testing will likely flush out problems that were not obvious in the confined settings of the developers’ sites; for example, it should help to resolve system and package dependencies, package interdependencies, and other details not addressed by this TDS version.

After most of the dust has settled, hopefully even conservative \TeX administrators will begin to adopt the TDS. Eventually, most \TeX sites will have adopted the common structure, and most packages will be readily available in TDS-compliant form.

We believe that this process will occur relatively quickly. The TDS committee spans a wide range of interests in the \TeX community. Consequently, we believe that most of the key issues involved in defining a workable TDS definition have been covered, often in detail. \TeX developers have been consulted about implementation issues, and have been trying out the TDS arrangement. Thus, we hope for few surprises as implementations mature.

Finally, there are several (current or prospective) publishers of \TeX CD-ROMs. These publishers are highly motivated to work out details of TDS implementation, and their products will provide inexpensive and convenient ways for experimentally-minded \TeX administrators to experiment with the TDS.

B.2 More on subdirectory searching

Recursive subdirectory searching is the ability to specify a search not only of a specified directory $\langle d \rangle$, but recursively of all directories below $\langle d \rangle$.

Since the TDS specifies precise locations for most files, with no extra levels of subdirectories allowed, true recursive searching is not actually required for a TDS-compliant implementation. We do, however, strongly recommend recursive searching as the most user-friendly and natural approach to the problem, rather than convoluted methods to specify paths without recursion.

This feature is already supported by many implementations of \TeX and companion utilities, for example DECUS \TeX for VMS, Dvips(k), em \TeX (and its drivers), PubliC \TeX , Web2C, Xdvi(k), and Y&Y \TeX . The Kpathsea library is a reusable implementation of subdirectory searching for \TeX , used in a number of the above programs.

Even if your \TeX implementation does not directly support subdirectory searching, you may find it useful to adopt the structure if you do not use many fonts or packages. For instance, if you only use Computer Modern and AMS fonts, it would be feasible to store them in the TDS layout and list the directories individually in configuration files or environment variables.

The TWG recognizes that subdirectory searching places an extra burden on the system and may be the source of performance bottlenecks, particularly on slower machines. Nevertheless, we feel that subdirectory searching is imperative for a well-organized TDS, for the reasons stated in Section 2.1. Implementors are encouraged to provide enhancements to the basic principle of subdirectory searching to avoid performance problems, e.g., the use of a filename cache (this can be as simple as a recursive directory listing) that is consulted before disk searching begins. If a match is found in the database, subdirectory searching is not required, and performance is thus independent of the number of subdirectories present on the system.

Different implementations specify subdirectory searching differently. In the interest of typographic clarity, the examples here do not use the $\langle replaceable \rangle$ font.

Dvips: via a separate `TEXFONTS_SUBDIR` environment variable.

em \TeX : `t:\subdir!!`; `t:\subdir!` for a single level of searching.

Kpathsea: `texmf/subdir//`

VMS: `texmf:[subdir...]`

Xdvi (patchlevel 20): `texmf/subdir/**`; `texmf/subdir/*` for a single level of searching. Version 20.50 and above support the `//` notation.

Y&Y \TeX : `t:/subdir//` or `t:\subdir\.`

B.3 Example implementation-specific trees

The TDS cannot specify a precise location for implementation-specific files, such as `texmf/ini`, because a site may have multiple \TeX implementations.

Nevertheless, for informative purposes, we provide here the default locations for some implementations. Please contact us with additions or corrections. These paths are not definitive, may not match anything at your site, and may change without warning.

We recommend all implementations have default search paths that start with the current directory (e.g., `'.'`). Allowing users to include the parent directory (e.g., `'..'`) is also helpful.

B.3.1 AmiWeb2c 2.0

(Email `scherer@physik.rwth-aachen.de` to contact the maintainer of this implementation.)

AmiWeb2c 2 is compatible with Web2c 7 to the greatest possible extent, so only the very few differences are described in this section. Detailed information about the basic concepts is given in the section for Web2c 7 below.

Thanks to the SELFAUTO mechanism of Kpathsea 3.0 no specific location for the installation of AmiWeb2c is required as long as the general structure of the distribution is preserved.

In addition to Kpathsea's // notation recursive path search may also be started by $\langle DEVICE \rangle:/$, e.g., `TeXMF:/` will scan this specific device completely.

Binaries coming with the AmiWeb2c distribution are installed in the directory `bin/amiweb2c/` outside the common TDS tree `share/texmf/`. In addition to the set of AmiWeb2c binaries you will find two subdirectories `local/` and `pastex/` with auxiliary programs.

A stripped version of the PasTeX system (used by kind permission of Georg Heßmann) is coming with AmiWeb2c, pre-installed in its own `share/texmf/amiweb2c/pastex/` directory. If you want to use PasTeX you have to assign the name `TeX:` to this place.

Documentation files in AmigaGuide format should be stored at `doc/guide/` similar to `doc/info/`.

B.3.2 Public DECUS TeX

If another VMS implementation besides Public DECUS TeX appears, the top level implementation directory name will be modified to something more specific (e.g., `vms-decus`).

```
texmf/  
  vms/          VMS implementation specific files  
    exe/        end-user commands  
      common/   command procedures, command definition files, etc.  
      axp/       binary executables for Alpha AXP  
      vax/       binary executables for VAX  
  formats/     pool files, formats, bases  
  help/        VMS help library, and miscellaneous help sources  
  mgr/         command procedures, programs, docs, etc., for system management
```

B.3.3 Web2c 7

All implementation-dependent TeX system files (`.pool`, `.fmt`, `.base`, `.mem`) are stored by default directly in `texmf/web2c`. The configuration file `texmf.cnf` and various subsidiary `MakeTeX...` scripts used as subroutines are also stored there.

Non-TeX specific files are stored following the GNU coding standards. Given a root directory $\langle prefix \rangle$ (`/usr/local` by default), we have default locations as follows:

<i><prefix>/</i>	installation root (<i>/usr/local</i> by default)
<i>bin/</i>	executables
<i>man/</i>	man pages
<i>info/</i>	info files
<i>lib/</i>	libraries (<i>libkpathsea.*</i>)
<i>share/</i>	architecture-independent files
<i>texmf/</i>	TDS root
<i>web2c/</i>	implementation-dependent files (<i>.pool</i> , <i>.fmt</i> , <i>texmf.cnf</i> , etc.)

See <http://www.gnu.org/prep/standards.toc.html> for the rationale behind and descriptions of this arrangement. A site may of course override these defaults; for example, it may put everything under a single directory such as */usr/local/texmf*.

C Is there a better way?

Defining the TDS required many compromises. Both the overall structure and the details of the individual directories were arrived at by finding common ground among many opinions. The driving forces were feasibility (in terms of what could technically be done and what could reasonably be expected from developers) and regularity (files grouped together in an arrangement that “made sense”).

Some interesting ideas could not be applied due to implementations lacking the necessary support:

- Path searching control at the \TeX level. If documents could restrict subdirectory searching to a subdirectory via some portable syntax in file names, restrictions on uniqueness of filenames could be relaxed considerably (with the cooperation of the formats), and the \TeX search path would not need to depend on the format.
- Multiple logical *texmf* trees. For example, a site might have one (read-only) location for stable files, and a different (writable) location for dynamically-created fonts or other files. It would be reasonable for two such trees to be logically merged when searching. See Michael Downes’ article in the references for how this can work in practice with Web2C.

C.1 Macro structure

The TWG settled on the *<format>/<package>* arrangement after long discussion about how best to arrange the files.

The primary alternative to this arrangement was a scheme which reversed the order of these directories: *<package>/<format>*. This reversed arrangement has a strong appeal: it keeps all of the files related to a particular package in a single place. The arrangement actually adopted tends to spread files out into two or three places (macros, documentation, and fonts, for example, are spread into different sections of the tree right at the top level).

Nevertheless, the *<format>/<package>* structure won for a couple of reasons:

- It is closer to current practice; in fact, several members of the TWG have already implemented the TDS hierarchy. The alternative is not in use at any known site, and the TWG felt it wrong to mandate something with which there is no practical experience.

- The alternative arrangement increases the number of top-level directories, so the files that must be found using subdirectory searching are spread out in a wide, shallow tree. This could have a profound impact on the efficiency of subdirectory searching.

C.2 Font structure

The TWG struggled more with the font directory structure than anything else. This is not surprising; the need to use the proliferation of PostScript fonts with $\text{T}_{\text{E}}\text{X}$ is what made the previous arrangement with all files in a single directory untenable, and therefore what initiated the TDS effort.

C.2.1 Font file type location

We considered the supplier-first arrangement in use at many sites:

```
texmf/fonts/<supplier>/<typeface>/<type>/
```

This improves the maintainability of the font tree, since all files comprising a given typeface are in one place, but unless all the programs that search this tree employ some form of caching, there are serious performance concerns. For example, in order to find a TFM file, the simplest implementation would require $\text{T}_{\text{E}}\text{X}$ to search through all the directories that contain PK files in all modes and at all resolutions.

In the end, a poll of developers revealed considerable resistance to implementing sufficient caching mechanisms, so this arrangement was abandoned. The TDS arrangement allows the search tree to be restricted to the correct type of file, at least. Concerns about efficiency remain, but there seems to be no more we can do without abandoning subdirectory searching entirely.

We also considered segregating all font-related files strictly by file type, so that METAFONT sources would be in a directory `texmf/fonts/mf`, property list files in `texmf/fonts/pl`, the various forms of Type 1 fonts separated, and so on. Although more blindly consistent, we felt that the drawback of more complicated path constructions outweighed this. The TDS merges file types (`mf` and `pl` under `source`, `pfa` and `pfb` and `gsf` under `type1`) where we felt this was beneficial.

C.2.2 Mode and resolution location

We considered having the `mode` at the bottom of the font tree:

```
texmf/fonts/pk/<supplier>/<typeface>/<mode>/<dpi>/
```

In this case, however, it is difficult to limit subdirectory searching to the mode required for a particular device.

We then considered moving the `dpi`*<nnn>* up to below the mode:

```
texmf/fonts/pk/<mode>/<dpi>/<supplier>/<typeface>/
```

But then it is not feasible to omit the `dpi`*<nnn>* level altogether on systems which can and do choose to use long filenames.

C.2.3 Modeless bitmaps

The TDS specifies using a single directory `modeless/` as the mode name for those utilities which generate bitmaps, e.g., `texmf/fonts/modeless/times/`. This has the considerable advantage of not requiring each such directory name to be listed in a search path.

An alternative was to use the utility name below which all such directories could be gathered. That has the advantage of separating, say, `gsftopk`-generated bitmaps from `ps2pk`-generated ones. However, we decided this was not necessary; most sites will use only one program for the purpose. Also, PK and GF fonts generally identify their creator in the font comment following the PK_ID byte.

We are making an implicit assumption that METAFONT is the only program producing mode-dependent bitmaps. If this becomes false we could add an abbreviation for the program to mode names, as in `mfcx` vs. `xyzcx` for a hypothetical program `XYZ`, or we could at that time add an additional program name level uniformly to the tree. It seemed more important to concisely represent the current situation than to worry about hypothetical possibilities that may never happen.

C.3 Documentation structure

We considered placing additional documentation files in the same directory as the source files for the packages, but we felt that users should be able to find documentation separately from sources, since most users have no interest in sources.

We hope that a separate, but parallel, structure for documentation would (1) keep the documentation together and (2) make it as straightforward as possible for users to find the particular documentation they were after.

D Related references

This appendix gives pointers to related files and other documents. For CTAN references, we use <http://www.ctan.org> as the top-level domain only to make the links be live in this document. See <http://www.ctan.org/tex-archive/CTAN.sites> for a complete list of CTAN sites; there are mirrors worldwide.

- This document, in many formats (tex, dvi, info, pdf):
<http://tug.org/tds/>
- The TDS mailing list archives:
<http://tug.org/mail-archives/twg-tds/>
- The level 0 DVI driver standard:
<http://www.ctan.org/tex-archive/dviware/driv-standard/level-0/>
- *Filenames for T_EX fonts*, with lists of recommended supplier and typeface names:
<http://tug.org/fontname/>
- ISO-9660 CD-ROM file system standard:
<http://www.iso.ch/cate/cat.html>
- *Components of T_EX*, a paper by Joachim Schrod:
<http://www.ctan.org/tex-archive/documentation/components-of-TeX/>

- *Managing Multiple TDS trees*, an article by Michael Downes:
<http://tug.org/TUGboat/Articles/tb22-3/tb72downes.pdf>
- A complete set of METAFONT modes:
<http://www.ctan.org/tex-archive/fonts/modes/modes.mf>
- A large collection of BIBTEX databases and styles:
<ftp://ftp.math.utah.edu/pub/tex/bib/>

E Contributors

The TWG has had no physical meetings; electronic mail was the communication medium.

Sebastian Rahtz is the T_EX Users Group Technical Council liaison. Norman Walsh was the original committee chair. Karl Berry is the current editor.

The list of contributors has grown too large to fairly include, as some would surely be inadvertently omitted. Please consider the archives of the `tds@tug.org` and `tex-live@tug.org` mailing lists as the record of contributions.