Fonts in XFree86

Juliusz Chroboczek, <jch@xfree86.org>

21 January 2001

1. Introduction

This document describes the support for fonts in XFree86. Section *Installing fonts* (section 2., page 1) is aimed at the casual user wishing to install fonts in the X server; the rest of the document describes the font support in more detail.

We only describe font support within the core X protocol. Issues relating to fonts within the RENDER extension, the GLX (OpenGL) extension or the PEX extension are outside the scope of this document.

We assume some familiarity with digital fonts. If anything is not clear to you, please consult Appendix *Background* (section 6., page 14) at the end of this document for background information.

2. Installing fonts

Installing fonts in XFree86 is a two step process. First, you need to create a *font directory* that contains all the relevant font files as well as some index files. You then need to inform the X server of the existence of this new directory by including it in the *font path*.

2.1 Installing bitmap fonts

The XFree86 server can use bitmap fonts in both the cross-platform BDF format and the somewhat more efficient binary PCF format. (XFree86 also supports the obsolete SNF format.)

Bitmap fonts are normally distributed in the BDF format. Before installing such fonts, it is desirable (but not absolutely necessary) to convert the font files to the PCF format. This is done by using the command '`bdftopcf`', *e.g.*

```
$ bdftopcf courier12.bdf
```

You may then want to compress the resulting PCF font files:

```
$ gzip courier12.pcf
```

After the fonts have been converted, you should copy all the font files that you wish to make available into a arbitrary directory, say '`/usr/local/share/fonts/bitmap/`'. You should then create the index file '`fonts.dir`' by running the command '`mkfontdir`' (please see the `mkfontdir`(1) manual page for more information):

```
$ mkdir /usr/local/share/fonts/bitmap
$ cp *.pcf.gz /usr/local/share/fonts/bitmap
$ cd /usr/local/share/fonts/bitmap
$ mkfontdir
```

All that remains is to tell the X server about the existence of the new font directory; see Section *Setting the server font path* (section 2.4, page 3).

2.2 Installing scalable fonts

The XFree86 server supports scalable fonts in four formats: Type 1, Speedo, TrueType and CID-Font. This section only applies to the former three; for information on CIDFonts, please see Section *Installing CIDFonts* (section 2.3, page 2) later in this document.

Installing scalable fonts is very similar to installing bitmap fonts: you create a directory with the font files, and run '`mkfontdir`' to create an index file called '`fonts.dir`'.

There is, however, a big difference: '`mkfontdir`' cannot automatically recognise scalable font files. For that reason, you must first index all the font files in a file called '`fonts.scale`'. This file has the same format as a '`fonts.dir`' file, and typically looks as follows:

```
4
cour.pfa -adobe-courier-medium-r-normal-0-0-0-0-p-0-iso8859-1
cour.pfa -adobe-courier-medium-r-normal-0-0-0-0-p-0-iso8859-2
couri.pfa -adobe-courier-medium-i-normal-0-0-0-0-p-0-iso8859-1
couri.pfa -adobe-courier-medium-i-normal-0-0-0-0-p-0-iso8859-2
```

The first line indicates the number of entries in the file. Each line after the first consists of two fields separated by a space; the first field is the name of the font file, and the second one is the name under which the font will appear to the server. This name should obey the X Logical Font Description conventions (see Section *The X Logical Font Description* (section 6.2, page 14)). The format of this file is fully described in the `mkfontdir`(1) manual page.

Note that multiple lines may point at the same font file. This is most commonly done in order to make a single font available under multiple encodings; please see Section *Fonts and internationalisation* (section 4., page 7).

While it is possible to create the '`fonts.scale`' file by hand, it is simpler and more convenient to have it generated automatically. Utilities to perform this task are available, but are not currently included with XFree86. For Type 1 fonts, you may use a utility called '`type1inst`' which is available from standard Free Software repositories <URL:http://www.ibiblio.org/pub/Linux/X11/xutils/> throughout the world.

For TrueType fonts, you may use '`ttmkfdir`', available from Joerg Pommnitz's xfsft page <URL:http://www.joerg-pommnitz.de/TrueType/xfsft.html>.

After the '`fonts.scale`' is created, you may run '`mkfontdir`' as above; this time, however, you need to create an index of encoding files called '`encodings.dir`' in addition to the '`fonts.dir`' file. This is done by using '`mkfontdir`' with the '`-e`' flag:

```
$ cd /usr/local/share/fonts/Type1
$ mkfontdir -e /usr/X11R6/lib/font/encodings
```

For more information, please see the `mkfontdir`(1) manual page and Section *Fonts and internationalisation* (section 4., page 7) later in this document.

2.3 Installing CID-keyed fonts

The CID-keyed font format was designed by Adobe Systems for fonts with large character sets. A CID-keyed font, or CIDFont for short, contains a collection of glyphs indexed by *character ID* (CID).

Adobe make some sample CIDFonts and a complete set of CMaps available from O'Reilly's FTP site <URL:ftp://ftp.oreilly.com/pub/examples/nutshell/cjkv/adobe/>.

In order to map such glyphs to meaningful indices, Adobe provide a set of *CMap* files. The

PostScript name of a font generated from a CIDFont consists of the name of the CIDFont and the name of the CMap separated by two dashes. For example, the font generated from the CIDFont '`Munhwa-Regular`' using the CMap '`UniKS-UCS2-H`' is called

```
Munhwa-Regular--UniKS-UCS2-H
```

The CIDFont support in XFree86 requires a very rigid directory structure. The main directory must be called '`CID`' (its location defaults to '`/usr/X11R6/lib/X11/fonts/CID`' but it may be located anywhere), and it should contain a subdirectory for every CID collection. Every subdirectory *must* contain subdirectories called `CIDFont` (containing the actual CIDFont files), `CMap` (containing all the needed CMaps), `AFM` (containing the font metric files) and `CFM` (initially empty). For example, in the case of the font `Munhwa-Regular` that uses the CID collection `Adobe-Korea1-0`, the directory structure should be as follows:

```
CID/Adobe-Korea1/CIDFont/Munhwa-Regular
CID/Adobe-Korea1/CMap/UniKS-UCS2-H
CID/Adobe-Korea1/AFM/Munhwa-Regular.afm
CID/Adobe-Korea1/CFM/
CID/fonts.dir
CID/fonts.scale
```

After creating this directory structure and copying the relevant files, you should create a <'tt/fonts.scale/' file. This file has the same format as in the case of (non-CID) scalable fonts, except that its first column contains PostScript font names with the extension '`.cid`' appended rather than actual filenames:

```
1
Adobe-Korea1/Munhwa-Regular--UniKS-UCS2-H.cid \
  -adobe-munhwa-medium-r-normal--0-0-0-0-p-0-iso10646-1
```

(both names on the same line). As above, running '`mkfontdir`' creates the '`fonts.dir`' file:

```
$ cd /usr/local/share/fonts/CID
$ mkfontdir
```

Finally, you should create the font metrics summary files in the directory '`CFM`' by running the command '`mkcfm`':

```
$ mkcfm /usr/local/share/fonts/CID
```

If no CFM files are available, the server will still be able to use the CID fonts but querying them will take a long time. You should run '`mkcfm`' again whenever a change is made to any of the CID-keyed fonts, or when the CID-keyed fonts are copied to a machine with a different architecture.

2.4 Setting the server's font path

The list of directories where the server looks for fonts is known as the *font path*. Informing the server of the existence of a new font directory consists in putting it on the font path.

The font path is an ordered list; if a client's request matches multiple fonts, the first one in the font path is the one that gets used. When matching fonts, the server makes two passes over the font path: during the first pass, it searches for an exact match; during the second, it searches for fonts suitable for scaling.

For best results, scalable fonts should appear in the font path before the bitmap fonts; this way, the server will prefer bitmap fonts to scalable fonts when an exact match is possible, but will avoid scaling bitmap fonts when a scalable font can be used. (The '`:unscaled`' hack, while still supported, should no longer be necessary in XFree86 4.0 and later.)

You may check the font path of the running server by typing the command

```
$ xset q
```

2.4.1 Temporary modification of the font path

The 'xset' utility may be used to modify the font path for the current session. The font path is set with the command `xset fp`; a new element is added to the front with `xset +fp`, and added to the end with `xset fp+`. For example,

```
$ xset +fp /usr/local/fonts/Type1
$ xset fp+ /usr/local/fonts/bitmap
```

Conversely, an element may be removed from the front of the font path with '`xset -fp`', and removed from the end with '`xset fp-`'.

For more information, please consult the `xset`(1) manual page.

2.4.2 Permanent modification of the font path

The default font path (the one used just after server startup) is specified in the X server's '`XF86Config`' file. It is computed by appending all the directories mentioned in the '`FontPath`' entries of the '`Files`' section in the order in which they appear.

```
FontPath "/usr/local/fonts/Type1"
...
FontPath "/usr/local/fonts/bitmap"
```

For more information, please consult the '`XF86Config`'(5) manual page.

2.5 Troubleshooting

If you seem to be unable to use some of the fonts you have installed, the first thing to check is that the '`fonts.dir`' files are correct and that they are readable by the server. If this doesn't help, it is quite possible that you are trying to use a font in a format that is not supported by your server.

XFree86 supports the BDF, PCF, SNF, Type 1, Speedo, TrueType and CIDFont font formats. However, not all XFree86 servers come with all the font backends configured in.

On most platforms, the XFree86 servers are *modular*: the font backends are included in modules that are loaded at runtime. The modules to be loaded are specified in the '`XF86Config`' file using the '`Load`' directive:

```
Load "type1"
```

If you have trouble installing fonts in a specific format, you may want to check the server's log file in order to see whether the relevant modules are properly loaded. The list of font modules distributed with XFree86 is as follows:

- "`bitmap`": bitmap fonts ('`*.bdf`', '`*.pcf`' and '`*.snf`');
- "`type1`": Type 1 fonts ('`*.pfa`' and '`*.pfb`') and CIDFonts;
- "`speedo`": Bitstream Speedo fonts ('`*.spd`');
- "`freetype`": TrueType fonts ('`*.ttf`' and '`*.ttc`');
- "`xtt`": alternate TrueType backend ('`*.ttf`' and '`*.ttc`').

Please note that the argument of the '`Load`' directive is case-sensitive.

3. Fonts included with XFree86

3.1 Standard bitmap fonts

The Sample Implementation of X11 comes with a large number of bitmap fonts, including the '`fixed`' family, and bitmap versions of Courier, Times and Helvetica. In the SI, these fonts are provided in the ISO 8859-1 encoding (ISO Latin Western-European).

In XFree86, a number of these fonts are provided in Unicode-encoded font files instead. At build time, these fonts are split into font files encoded according to legacy encodings, a process which enables us to provide the standard fonts in a number of regional encodings with no duplication of work.

For example, the font file

```
/usr/X11R6/lib/X11/fonts/misc/6x13.bdf
```

with XLFD

```
-misc-fixed-medium-r-semicondensed--13-120-75-75-c-60-iso10646-1
```

is a Unicode-encoded version of the standard '`fixed`' font with added support for the Latin, Greek, Cyrillic, Georgian, Armenian, IPA and other scripts plus numerous technical symbols. It contains over 2800 glyphs, covering all characters of ISO 8859 parts 1-5, 7-10, 13-15, as well as all European IBM and Microsoft code pages, KOI8, WGL4, and the repertoires of many other character sets.

This font is used at build time for generating the font files

```
6x13-ISO8859-1.bdf
6x13-ISO8859-2.bdf
6x13-ISO8859-3.bdf
6x13-ISO8859-4.bdf
6x13-ISO8859-5.bdf
6x13-ISO8859-7.bdf
6x13-ISO8859-8.bdf
6x13-ISO8859-9.bdf
6x13-ISO8859-10.bdf
6x13-ISO8859-13.bdf
6x13-ISO8859-15.bdf
6x13-KOI8-R.bdf
```

with respective XLFDs

```
-misc-fixed-medium-r-normal--13-120-75-75-c-60-iso8859-1
-misc-fixed-medium-r-normal--13-120-75-75-c-60-iso8859-2
-misc-fixed-medium-r-normal--13-120-75-75-c-60-iso8859-3
-misc-fixed-medium-r-normal--13-120-75-75-c-60-iso8859-4
-misc-fixed-medium-r-normal--13-120-75-75-c-60-iso8859-5
-misc-fixed-medium-r-normal--13-120-75-75-c-60-iso8859-7
-misc-fixed-medium-r-normal--13-120-75-75-c-60-iso8859-8
-misc-fixed-medium-r-normal--13-120-75-75-c-60-iso8859-9
-misc-fixed-medium-r-normal--13-120-75-75-c-60-iso8859-10
-misc-fixed-medium-r-normal--13-120-75-75-c-60-iso8859-13
-misc-fixed-medium-r-normal--13-120-75-75-c-60-iso8859-15
-misc-fixed-medium-r-normal--13-120-75-75-c-60-koi8-r
```

The standard short name '`fixed`' is normally an alias for

```
-misc-fixed-medium-r-normal--13-120-75-75-c-60-iso8859-1
```

(The conversion of the standard fonts to Unicode was mainly performed by Markus Kuhn. Markus is a man of taste, which makes his use of Perl in the conversion process somewhat surprising.)

3.2 The ClearlyU Unicode font family

The ClearlyU family of fonts provides a set of 12 pt, 100 dpi proportional fonts with many of the glyphs needed for Unicode text. Together, the fonts contain approximately 7500 glyphs.

The main ClearlyU font has the XLFD

```
-mutt-clearlyu-medium-r-normal--17-120-100-100-p-101-iso10646-1
```

and resides in the font file

```
/usr/X11R6/lib/X11/fonts/misc/cu12.pcf.gz
```

Additional ClearlyU fonts include

```
-mutt-clearlyu alternate glyphs-medium-r-normal--17-120-100-100-p-91-iso10646-1
-mutt-clearlyu pua-medium-r-normal--17-120-100-100-p-111-iso10646-1
-mutt-clearlyu arabic extra-medium-r-normal--17-120-100-100-p-103-fontspecific-0
-mutt-clearlyu ligature-medium-r-normal--17-120-100-100-p-141-fontspecific-0
```

The *Alternate Glyphs* font contains additional glyph shapes that are needed for certain languages. A second alternate glyph font will be provided later for cases where a character has more than one commonly used alternate shape (*e.g.* the Urdu heh).

The *PUA* font contains extra glyphs that are useful for certain rendering purposes.

The *Arabic Extra* font contains the glyphs necessary for characters that don't have all of their possible shapes encoded in ISO 10646. The glyphs are roughly ordered according to the order of the characters in the ISO 10646 standard.

The *Ligature* font contains ligatures for various scripts that may be useful for improved presentation of text.

(The ClearlyU family was designed by Mark Leisher. Mark's usage of the foundry name *mutt* predates the mailer of the same name, but he won't say more.)

3.3 Standard scalable fonts

XFree86 includes all the scalable fonts distributed with X11R6.

3.3.1 Standard Type 1 fonts

The IBM Courier set of fonts cover ISO 8859-1 and ISO 8859-2 as well as Adobe Standard Encoding. These fonts have XLFD

```
-adobe-courier-medium-*-*--0-0-0-0-m-0-*-*
```

and reside in the font files

```
/usr/X11R6/lib/X11/fonts/Type1/cour*.pfa
```

The Adobe Utopia set of fonts only cover ISO 8859-1 as well as Adobe Standard Encoding. These fonts have XLFD

```
-adobe-utopia-*-*-normal--0-0-0-0-p-0-iso8859-1
```

and reside in the font files

```
/usr/X11R6/lib/X11/fonts/Type1/UT*.pfa
```

Finally, XFree86 also comes with Type 1 versions of Bitstream Courier and Charter. These fonts have XLFD

```
-bitstream-courier-*-*-normal--0-0-0-0-m-0-iso8859-1
-bitstream-charter-*-*-normal--0-0-0-0-p-0-iso8859-1
```

and reside in the font files

```
/usr/X11R6/lib/X11/fonts/Type1/c*bt_.pfb
```

3.3.2 Standard Speedo fonts

XFree86 includes Speedo versions of the Bitstream Courier and Charter fonts. In order to use these fonts, you should ensure that your X server is loading the '`Speedo`' font backend; see Section *Troubleshooting* (section 2.5, page 4).

These fonts cover all of ISO 8859-1 and almost all of ISO 8859-2. They have XLFD name

```
-bitstream-courier-*-*-normal--0-0-0-0-m-0-*-*
-bitstream-charter-*-*-normal--0-0-0-0-p-0-*-*
```

and reside in the font files

```
/usr/X11R6/lib/X11/fonts/Speedo/font*.spd
```

3.4 The Bigelow & Holmes Lucidux family

XFree86 includes the *Lucidux* family of Type 1 fonts. This family consists of the fonts *Lucidux Serif*, with XLFD

```
-b&h-lucidux serif-medium-*-normal--*-*-*-*-p-*-*-*
```

Lucidux Sans, with XLFD

```
-b&h-lucidux sans-medium-*-normal--*-*-*-*-p-*-*-*
```

and *Lucidux Mono*, with XLFD

```
-b&h-lucidux mono-medium-*-normal--*-*-*-*-m-*-*-*
```

Each of these fonts currently comes in Roman and oblique variants (bold variants will be included in a future release) and has 337 glyphs covering the basic ASCII Unicode range, the Latin 1 range, as well as the *Extended Latin* range. In particular, these fonts include all the glyphs needed for ISO 8859 parts 1, 2, 3, 4, 9 and 15.

The Lucidux fonts are original designs by Charles Bigelow and Kris Holmes. Lucidux fonts include seriffed, sans serif, and monospaced styles that share the same stem weight, x-height, capital height, ascent and descent. Lucidux fonts harmonise with Lucida (R) fonts of the same vertical proportions and weights. The character width metrics of Lucidux roman fonts match those of core fonts bundled with several window systems.

Each PFA file has a copy of the license terms in PS comment lines. The license terms are also included in the file '`COPYRIGHT.BH`' for convenience, as well as in the License document.

The design and font outlines were donated by Charles Bigelow and Kris Holmes from Bigelow and Holmes Inc., and the hinting was donated by Berthold Horn and Blenda Horn from Y&Y, Inc. For more information, please contact <design@bigelowandholmes.com> or <sales@yandy.com>, or consult Y&Y's web site <URL:`http://www.yandy.com`>.

4. Fonts and internationalisation

The scalable font backends (Type 1, Speedo and TrueType) can now automatically re-encode fonts to the encoding specified in the XLFD in `fonts.dir`. For example, a `fonts.dir` file can

contain entries for the Type 1 Courier font such as

```
cour.pfa -adobe-courier-medium-r-normal--0-0-0-0-m-0-iso8859-1
cour.pfa -adobe-courier-medium-r-normal--0-0-0-0-m-0-iso8859-2
```

which will lead to the font being recoded to ISO 8859-1 and ISO 8859-2 respectively.

4.1 The fontenc layer

Three of the scalable backends (Type 1, Speedo, and the *FreeType* TrueType backend) use a common *fontenc* layer for font re-encoding. This allows these backends to share their encoding data, and allows simple configuration of new locales independently of font type.

Please note: the X-TrueType (X-TT) backend does not use the *fontenc* layer, but instead uses its own method for font reencoding. If you are only interested in X-TT you may want to skip to Section *Using Symbol Fonts* (section 4.5, page 12), as the intervening information does not apply to X-TT. X-TT itself is described in more detail in Section *X-TrueType* (section 5.2, page 13).

In the *fontenc* layer, an encoding is defined by a name (such as `iso8859-1`), possibly a number of aliases (alternate names), and an ordered collection of mappings. A mapping defines the way the encoding can be mapped into one of the *target encodings* known to *fontenc*; currently, these consist of Unicode, Adobe glyph names, and arbitrary TrueType ''cmap''s.

A number of encodings are hardwired into *fontenc*, and are therefore always available; the hard-coded encodings cannot easily be redefined. These include:

- `iso10646-1`: Unicode;
- `iso8859-1`: ISO Latin-1 (Western Europe);
- `iso8859-2`: ISO Latin-2 (Eastern Europe);
- `iso8859-3`: ISO Latin-3 (Southern Europe);
- `iso8859-4`: ISO Latin-4 (Northern Europe);
- `iso8859-5`: ISO Cyrillic;
- `iso8859-6`: ISO Arabic;
- `iso8859-7`: ISO Greek;
- `iso8859-8`: ISO Hebrew;
- `iso8859-9`: ISO Latin-5 (Turkish);
- `iso8859-10`: ISO Latin-6 (Nordic);
- `iso8859-15`: ISO Latin-9, or Latin-0 (Revised Western-European);
- `koi8-r`: KOI8 Russian;
- `koi8-u`: KOI8 Ukrainian (see RFC 2319);
- `koi8-ru`: KOI8 Russian/Ukrainian
- `koi8-uni`: KOI8 ''Unified'' (Russian, Ukrainian, and Byelorussian);
- `koi8-e`: KOI8 ''European,'' ISO-IR-111, or ECMA-Cyrillic;
- `microsoft-symbol` and `apple-roman`: these are only likely to be useful with TrueType symbol fonts.

Additional encodings can be added by defining *encoding files*. When a font encoding is requested that the *fontenc* layer doesn't know about, the backend checks the directory in which the font file resides (not necessarily the directory with `fonts.dir`!) for a file named '`encodings.dir`'. If

found, this file is scanned for the requested encoding, and the relevant encoding definition file is read in. The 'mkfontdir' utility, when invoked with the '-e' option followed by the name of a directory containing encoding files, can be used to automatically build 'encodings.dir' files. See the mkfontdir(1) manual page for more details.

A number of encoding files for common encodings are included with XFree86. Information on writing new encoding files can be found in Section *Format of encodings directory files* (section 4.3, page 9) and *Format of encoding files* (section 4.4, page 10) later in this document.

4.2 Backend-specific notes about fontenc

4.2.1 Type 1

The Type 1 backend first searches for a mapping with a target of PostScript. If one is found, it is used. Otherwise, the backend searches for a mapping with target Unicode, which is then composed with a built-in table mapping codes to glyph names. Note that this table only covers part of the Unicode code points that have been assigned names by Adobe.

If neither a PostScript or Unicode mapping is found, the backend defaults to ISO 8859-1.

Specifying an encoding value of adobe-fontspecific disables the encoding mechanism. This is useful with symbol and incorrectly encoded fonts (see Section *Incorrectly encoded fonts* (section 4.6, page 12) below).

The Type 1 backend currently limits all encodings to 8-bit codes.

4.2.2 Speedo

The Speedo backend searches for a mapping with a target of Unicode, and uses it if found. If none is found, the backend defaults to ISO 8859-1.

The Speedo backend limits all encodings to 8-bit codes.

4.2.3 The FreeType TrueType backend

The TrueType backend scans the mappings in order. Mappings with a target of PostScript are ignored; mappings with a TrueType or Unicode target are checked against all the cmaps in the file. The first applicable mapping is used.

If you are writing an encoding file to be used with the TrueType backend, you should ensure that mappings are mentioned in decreasing order of preference.

4.3 Format of encoding directory files

In order to use a font in an encoding that the font backend does not know about, you need to have an 'encodings.dir' file in the same directory as the font file used. The 'encodings.dir' file has a similar format to 'fonts.dir'. Its first line specifies the number of encodings, while every successive line has two columns, the name of the encoding, and the name of the encoding file; this can be relative to the current directory, or absolute. Every encoding name should agree with the encoding name defined in the encoding file. For example,

```
3
mulearabic-0 /usr/X11R6/lib/X11/fonts/encodings/mulearabic-0.enc
mulearabic-1 /usr/X11R6/lib/X11/fonts/encodings/mulearabic-1.enc
mulearabic-2 /usr/X11R6/lib/X11/fonts/encodings/mulearabic-2.enc
```

The name of an encoding *must* be specified in the encoding file's 'STARTENCODING' or 'ALIAS' line. It is not enough to create an 'encodings.dir' entry.

If your platform supports it (it probably does), encoding files may be compressed or gzipped.

The 'encoding.dir' files are best maintained by the 'mkfontdir' utility. Please see the mkfontdir(1) manual page for more information.

4.4 Format of encoding files

The encoding files are ''free form,'' *i.e.* any string of whitespace is equivalent to a single space. Keywords are parsed in a non-case-sensitive manner, meaning that '`size`', '`SIZE`', and '`SiZE`' all parse as the same keyword; on the other hand, case is significant in glyph names.

Numbers can be written in decimal, as in '`256`', in hexadecimal, as in '`0x100`', or in octal, as in '`0400`'.

Comments are introduced by a hash sign '`#`'. A '`#`' may appear at any point in a line, and all characters following the '`#`' are ignored, up to the end of the line.

The encoding file starts with the definition of the name of the encoding, and possibly its alternate names (aliases):

```
STARTENCODING mulearabic-0
ALIAS arabic-0
ALIAS something-else
```

The name of the encoding and its aliases should be suitable for use in an XLFD font name, and therefore contain exactly one dash '`-`'.

The encoding file may then optionally declare the size of the encoding. For a linear encoding (such as ISO 8859-1), the SIZE line specifies the maximum code plus one:

```
SIZE 0x2B
```

For a matrix encoding, it should specify two numbers. The first is the number of the last row plus one, the other, the highest column number plus one. In the case of '`jisx0208.1990-0`' (JIS X 0208(1990), double-byte encoding, high bit clear), it should be

```
SIZE 0x75 0x80
```

In the case of a matrix encoding, a '`FIRSTINDEX`' line may be included to specify the minimum glyph index in an encoding. The keyword '`FIRSTINDEX`' is followed by two integers, the minimum row number followed by the minimum column number:

```
FIRSTINDEX 0x20 0x20
```

In the case of a linear encoding, a '`FIRSTINDEX`' line is not very useful. If for some reason however you chose to include on, it should be followed by a single integer.

Note that in most font backends inclusion of a '`FIRSTINDEX`' line has the side effect of disabling default glyph generation, and this keyword should therefore be avoided unless absolutely necessary.

Codes outside the region defined by the '`SIZE`' and '`FIRSTINDEX`' lines are understood to be undefined. Encodings default to linear encoding with a size of 256 (0x100). This means that you must declare the size of all 16 bit encodings.

What follows is one or more mapping sections. A mapping section starts with a '`STARTMAPPING`' line stating the target of the mapping. The target may be one of:

- Unicode (ISO 10646):

  ```
  STARTMAPPING unicode
  ```

- a given TrueType ''cmap'':

  ```
  STARTMAPPING cmap 3 1
  ```

- PostScript glyph names:

```
STARTMAPPING postscript
```

Every line in a mapping section maps one from the encoding being defined to the target of the mapping. In mappings with a Unicode or TrueType mapping, codes are mapped to codes:

```
0x21 0x0660
0x22 0x0661
...
```

As an abbreviation, it is possible to map a contiguous range of codes in a single line. A line consisting of three integers

```
<it/start/ <it/end/ <it/target/
```

is an abbreviation for the range of lines

start *target*

start+1 *target*+1

...

end *target+end-start*

For example, the line

```
0x2121 0x215F 0x8140
```

is an abbreviation for

```
0x2121 0x8140
0x2122 0x8141
...
0x215F 0x817E
```

Codes not listed are assumed to map through the identity (*i.e.* to the same numerical value). In order to override this default mapping, you may specify a range of codes to be undefined by using an 'UNDEFINE' line:

```
UNDEFINE 0x00 0x2A
```

or, for a single code,

```
UNDEFINE 0x1234
```

PostScript mappings are different. Every line in a PostScript mapping maps a code to a glyph name

```
0x41 A
0x42 B
...
```

and codes not explicitly listed are undefined.

A mapping section ends with an ENDMAPPING line

```
ENDMAPPING
```

After all the mappings have been defined, the file ends with an ENDENCODING line

```
ENDENCODING
```

In order to make future extensions to the format possible, lines starting with an unknown keyword are silently ignored, as are mapping sections with an unknown target.

4.5 Using symbol fonts

Type 1 symbol fonts should be installed using the `adobe-fontspecific` encoding.

In an ideal world, all TrueType symbol fonts would be installed using one of the `microsoft-symbol` and `apple-roman` encodings. A number of symbol fonts, however, are not marked as such; such fonts should be installed using `microsoft-cp1252`, or, for older fonts, `microsoft-win3.1`.

In order to guarantee consistent results (especially between Type 1 and TrueType versions of the same font), it is possible to define a special encoding for a given font. This has already been done for the `ZapfDingbats` font; see the file `encodings/adobe-dingbats.enc`.

4.6 Hints about using badly encoded fonts

A number of text fonts are incorrectly encoded. Incorrect encoding is sometimes done by design, in order to make a font for an exotic script appear like an ordinary Western text font. It is often the result of the font designer's laziness or incompetence; for some reason, most people seem to find it easier to invent idiosyncratic glyph names rather than follow the Adobe glyph list.

There are two ways of dealing with such fonts: using them with the encoding they were designed for, and creating an *ad hoc* encoding file.

4.6.1 Using fonts with the designer's encoding

In the case of Type 1 fonts, the font designer can specify a default encoding; this encoding is requested by using the '`adobe-fontspecific`' encoding in the XLFD name. Sometimes, the font designer omitted to specify a reasonable default encoding, in which case you should experiment with '`adobe-standard`', '`iso8859-1`', '`microsoft-cp1252`', and '`microsoft-win3.1`'. (The encoding '`microsoft-symbol`' doesn't make sense for Type 1 fonts).

TrueType fonts do not have a default encoding. However, most TrueType fonts are designed with either Microsoft or Apple platforms in mind, so one of '`microsoft-symbol`', '`microsoft-cp1252`', '`microsoft-win3.1`', or '`apple-roman`' should yield reasonable results.

4.6.2 Specifying an ad hoc encoding file

It is always possible to define an encoding file to put the glyphs in a font in any desired order. Again, see the '`encodings/adobe-dingbats.enc`' file to see how this is done.

4.6.3 Specifying font aliases

By following the directions above, you will find yourself with a number of fonts with unusual names --- with encodings such as '`adobe-fontspecific`', '`microsoft-win3.1`' *etc*. In order to use these fonts with standard applications, it may be useful to remap them to their proper names.

This is done by writing a '`fonts.alias`' file. The format of this file is similar to the format of the '`fonts.dir`' file, except that it maps XLFD names to XLFD names. A '`fonts.alias`' file might look as follows:

```
1
"-ogonki-alamakota-medium-r-normal--0-0-0-0-p-0-iso8859-2" \
  "-ogonki-alamakota-medium-r-normal--0-0-0-0-p-0-adobe-fontspecific"
```

(both XLFD names on a single line). The syntax of the '`fonts.alias`' file is precisely described in the `mkfontdir`(1) manual page.

5. Additional notes about TrueType support

This version of XFree86 comes with two TrueType backends, *FreeType* (module 'freetype', formerly known as *xfsft*) and *X-TrueType* (module 'xtt'). These two backends are *not* compatible: only one of them can be used at any one time.

In order to use the *FreeType* backend, please check that the 'Module' section of your 'XF86Config' file contains a line that reads

```
Load "freetype"
```

In order to use the *X-TrueType* backend, replace the line in your XF86Config file that loads the freetype module with a line that reads

```
Load "xtt"
```

Both TrueType backends delay glyph rasterisation up to the time at which a glyph is first used. For this reason, they only provide an approximate value for the ''average width'' font property.

Both backends also support an optimisation for character-cell fonts (fonts with all glyph metrics equal, or terminal fonts). A font with an XLFD specifying a character-cell spacing 'c', as in

```
-misc-mincho-medium-r-normal--0-0-0-0-c-0-jisx0208.1990-0
```

will not rasterise glyphs at metrics computation time, but instead trust the font really to be a character-cell font. You are encouraged to make use of this optimisation when useful, but be warned that not all monospaced fonts are character-cell fonts.

5.1 The FreeType TrueType backend

The *FreeType* backend (formerly *xfsft*) is a backend based on the FreeType library (see the FreeType web site <URL:http://www.freetype.org/>) and has support for the ''fontenc'' style of internationalisation (see Section *The fontenc layer* (section 4.1, page 8)). This backend supports TrueType Font files ('*.ttf') and TrueType Collections ('*.ttc').

In order to access the faces in a TrueType Collection file, the face number must be specified in the fonts.dir file before the filename within colons. For example,

```
:2:mincho.ttc -misc-mincho-medium-r-normal--0-0-0-0-c-0-jisx0208.1990-0
```

refers to face 2 in the 'mincho.ttc' TrueType Collection file.

The *FreeType* backend uses the *fontenc* layer in order to support recoding of fonts; this was described in Section *The fontenc layer* (section 4.1, page 8) and especially Section *FreeType-specific notes about fontenc* (section 4.2.3, page 9) earlier in this document.

5.2 The X-TrueType TrueType backend

The 'X-TrueType' backend is another backend based on the FreeType library. X-TrueType doesn't use the 'fontenc' layer for managing font encodings, but instead uses its own database of encodings. However, X-TrueType includes a large number of encodings, and any encoding you need is likely to be present in X-TrueType.

X-TrueType extends the 'fonts.dir' syntax with a number of options, known as 'TTCap'. A 'TTCap' entry follows the general syntax

```
:option=value:
```

and should be specified before the filename.

The most useful TTCap option is used to specify the face number to use with TTCs; this is the 'fn' TTCap option. For example, face 2 of font file 'mincho.ttc' is specified using:

```
:fn=2:mincho.ttc -misc-mincho-medium-r-normal--0-0-0-0-c-0-jisx0208.1990-0
```

More information on the TTCap syntax, and on X-TrueType in general, may be found on the X-TrueType home page <URL:http://x-tt.dsl.gr.jp/>.

6. Appendix: background and terminology

6.1 Characters and glyphs

A computer text-processing system inputs keystrokes and outputs *glyphs*, small pictures that are assembled on paper or on a computer screen. Keystrokes and glyphs do not, in general, coincide: for example, if the system does generate ligatures, then to the two keystrokes <f><i> will typically correspond a single glyph. Similarly, if the system shapes Arabic glyphs in a reasonable manner, then multiple different glyphs may correspond to a single keystroke.

The complex transformation rules from keystrokes to glyphs are usually factored into two simpler transformations, going through the intermediary of *characters*. You may want to think of characters as the basic unit of data that is stored *e.g.* in the buffer of your text editor. While the definition of a character is intrinsically application-specific, a number of standardised collections of characters have been defined.

A *coded character set* is a set of characters together with a mapping from integer codes --- known as *codepoints* --- to characters. Examples of coded character sets include US-ASCII, ISO 8859-1, KOI8-R, and JIS X 0208(1990).

A coded character set need not use 8 bit integers to index characters. Many early mainframes used 6 bit character sets, while 16 bit (or more) character sets are necessary for ideographic writing systems.

6.2 Font files, fonts, and XLFD

Traditionally, typographers speak about *typefaces* and *founts*. A typeface is a particular style or design, such as Times Italic, while a fount is a molten-lead incarnation of a given typeface at a given size.

Digital fonts come in *font files*. A font file contains all the information necessary for generating glyphs of a given typeface, and applications using font files may access glyph information in an arbitrary order.

Digital fonts may consist of bitmap data, in which case they are said to be *bitmap fonts*. They may also consist of a mathematical description of glyph shapes, in which case they are said to be *scalable fonts*. Common formats for scalable font files are *Type 1* (sometimes incorrectly called *ATM fonts* or *PostScript fonts*), *Speedo* and *TrueType*.

The glyph data in a digital font needs to be indexed somehow. How this is done depends on the font file format. In the case of Type 1 fonts, glyphs are identified by *glyph names*. In the case of TrueType fonts, glyphs are indexed by integers corresponding to one of a number of indexing schemes (usually Unicode --- see below).

The X11 system uses the data in font file to generate *font instances*, which are collections of glyphs at a given size indexed according to a given encoding.

X11 font instances are usually specified using a notation known as the *X Logical Font Description* (XLFD). An XLFD starts with a dash '-', and consists of fourteen fields separated by dashes, for example

```
-adobe-courier-medium-r-normal--0-0-0-0-m-0-iso8859-1
```

Or particular interest are the last two fields '`iso8859-1`', which specify the font instance's

encoding.

X11 font instances may also be specified by short name. Unlike an XLFD, a short name has no structure and is simply a conventional name for a font instance. Two short names are of particular interest, as they are handled specially by the server, and the server will not start if font instances with these names cannot be opened. These are '`fixed`', which specifies the fallback font to use when the requested font cannot be opened, and '`cursor`', which specifies the set of glyphs to be used by the mouse pointer.

Short names are usually implemented as aliases to XLFDs; the '`fixed`' and '`cursor`' aliases are defined in

```
/usr/X11R6/lib/X11/font/misc/fonts.alias
```

6.3 Unicode

Unicode (<URL:http://www.unicode.org>) is a coded character set with the goal of uniquely identifying all characters for all scripts, current and historical. While Unicode was explicitly not designed as a glyph encoding scheme, it is often possible to use it as such.

Unicode is an *open* character set, meaning that codepoint assignments may be added to Unicode at any time (once specified, though, an assignment can never be changed). For this reason, a Unicode font will be *sparse*, and only define glyphs for a subset of the character registry of Unicode.

The Unicode standard is defined in parallel with the international standard ISO 10646. Assignments in the two standards are always equivalent, and this document uses the terms *Unicode* and *ISO 10646* interchangeably.

When used in X11, Unicode-encoded fonts should have the last two fields of their XLFD set to '`iso10646-1`'.

7. References

XFree86 comes with extensive documentation in the form of manual pages and typeset documents. Before installing fonts, you really should read the `mkfontdir`(1) manual page; other manual pages of interest include `X`(1), `Xserver`(1), `xset`(1), `xlsfonts`(1) and `showfont`(1). In addition, you may want to read the X Logical Font Description document, by Jim Flowers, which is provided in the file '`xc/doc/xlfd.PS.Z`'.

The comp.fonts FAQ <URL:http://www.netmeg.net/faq/computers/fonts/>, which is unfortunately no longer being maintained, contains a wealth of information about digital fonts.

The xfsft home page <URL:http://www.dcs.ed.ac.uk/home/jec/programs/xfsft/> has been superseded by this document, and is now obsolete; you may however still find some of the information it contains useful. Joerg Pommnitz' xfsft page <URL:http://www.joerg-pommnitz.de/TrueType/xfsft.html> is the canonical source for the '`ttmkfdir`' utility.

The documentation of *X-TrueType* is available from the X-TrueType home page <URL:http://x-tt.dsl.gr.jp/>.

A number of East-Asian CIDFonts are available from O'Reilly's FTP site <URL:ftp://ftp.oreilly.com/pub/examples/nutshell/cjkv/adobe/>.

The Unicode consortium site <URL:http://www.unicode.org> may be of interest. But you are more likely to find what you need on Markus Kuhn's UTF-8 and Unicode FAQ <URL:http://www.cl.cam.ac.uk/~mgk25/unicode.html>.

The IANA RFC documents, available from a number of sites throughout the world, often provide interesting information about character set issues; my favourite is RFC 373.

CONTENTS

1. Introduction 1

2. Installing fonts 1
   - 2.1 Installing bitmap fonts 1
   - 2.2 Installing scalable fonts 2
   - 2.3 Installing CID-keyed fonts 2
   - 2.4 Setting the server's font path 3
   - 2.5 Troubleshooting 4

3. Fonts included with XFree86 4
   - 3.1 Standard bitmap fonts 5
   - 3.2 The ClearlyU Unicode font family 6
   - 3.3 Standard scalable fonts 6
   - 3.4 The Bigelow & Holmes Lucidux family 7

4. Fonts and internationalisation 7
   - 4.1 The fontenc layer 8
   - 4.2 Backend-specific notes about fontenc 9
   - 4.3 Format of encoding directory files 9
   - 4.4 Format of encoding files 10
   - 4.5 Using symbol fonts 12
   - 4.6 Hints about using badly encoded fonts 12

5. Additional notes about TrueType support 13
   - 5.1 The FreeType TrueType backend 13
   - 5.2 The X-TrueType TrueType backend 13

6. Appendix: background and terminology 14
   - 6.1 Characters and glyphs 14
   - 6.2 Font files, fonts, and XLFD 14
   - 6.3 Unicode 15

7. References 15

`$XFree86: xc/programs/Xserver/hw/xfree86/doc/sgml/fonts.sgml,v 1.11.2.1 2001/02/28 20:06:21 dawes Exp`