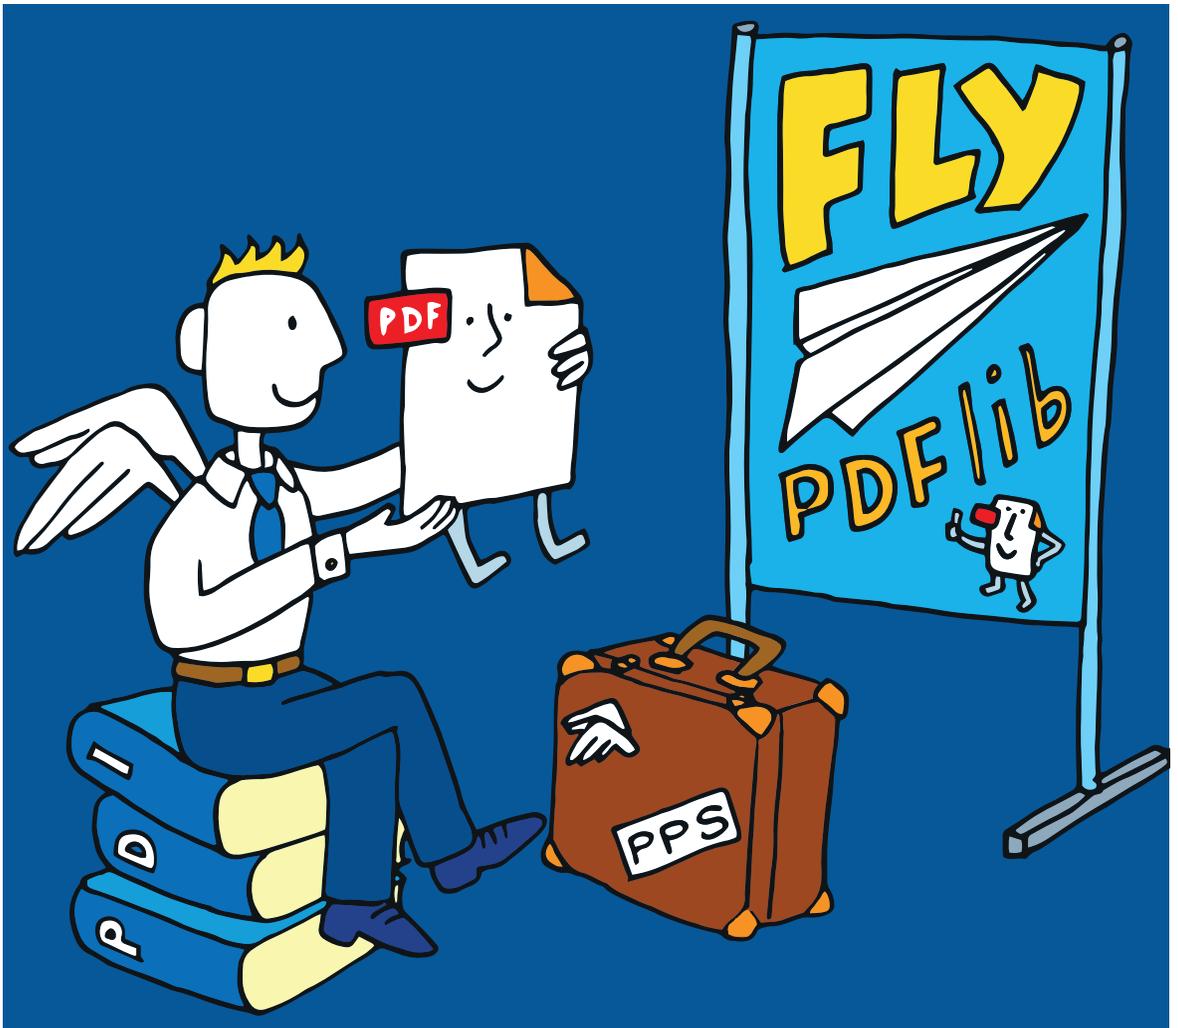


PDFlib, PDFlib+PDI, PPS

A library for generating PDF on the fly
PDFlib 10.0.1

Tutorial

For use with C, C++, Java, .NET, Objective-C, Perl, PHP, Python, RPG, Ruby



Copyright © 1997–2022 PDFlib GmbH. All rights reserved.
PDFlib users are granted permission to reproduce printed or digital copies of this manual for internal use.

PDFlib GmbH
Franziska-Bilek-Weg 9, 80339 München, Germany
www.pdflib.com
phone +49 • 89 • 452 33 84-0

sales@pdflib.com
support@pdflib.com (please include your license number)

This publication and the information herein is furnished as is, is subject to change without notice, and should not be construed as a commitment by PDFlib GmbH. PDFlib GmbH assumes no responsibility or liability for any errors or inaccuracies, makes no warranty of any kind (express, implied or statutory) with respect to this publication, and expressly disclaims any and all warranties of merchantability, fitness for particular purposes and noninfringement of third party rights.

PDFlib and the PDFlib logo are registered trademarks of PDFlib GmbH. PDFlib licensees are granted the right to use the PDFlib name and logo in their product documentation. However, this is not required.

PDFlib contains PANTONE® spot color data © Pantone LLC, 2019. PANTONE® colors displayed in the software application or in the user documentation may not match PANTONE-identified standards. Consult current PANTONE Color Publications for accurate color. PANTONE® and other Pantone, Inc. trademarks are the property of Pantone, Inc. © Pantone, LLC, 2019. Pantone, Inc. is the copyright owner of color data and/or software which are licensed to PDFlib GmbH to distribute for use only in combination with PDFlib Software. PANTONE Color Data and/or Software shall not be copied onto another disk or into memory unless as part of the execution of PDFlib Software.

PDFlib contains HKS® spot color data licensed from HKS Warenzeichenverband e.V.

PDFlib contains the following third-party components:
Adobe CMap resources, Copyright © 1990-2019 Adobe
AES, Arcfour and SHA algorithms, Copyright © 1995-1998 Eric Young
Brotli decompression code, Copyright © 2009, 2010, 2013-2016 by the Brotli Authors
Curl multiprotocol file transfer library, Copyright © 1996-2022 Daniel Stenberg
Expat XML parser, Copyright © 2001-2019 Expat maintainers
ICCLib, Copyright © 1997-2021 Graeme W. Gill
ICU International Components for Unicode, Copyright © 1991-2020 Unicode, Inc.
libjpeg, Copyright © 1991-2020, Thomas G. Lane, Guido Vollbeding
libpng, Copyright © 2006-2017 Glenn Randers-Pehrson, Copyright © 2018-2019 Cosmin Truta
Libtiff image library, Copyright © 1988-1997 Sam Leffler, Copyright © 1991-1997 Silicon Graphics, Inc.
MD5 message digest, Copyright © 1991-2, RSA Data Security, Inc.
NotoSans fonts, Copyright © 2012 Google Inc. All Rights Reserved.
OpenJPEG library, Copyright © 2002-2014, Université catholique de Louvain (UCL), Belgium
OpenSSL cryptographic library, Copyright © 1998-2018 The OpenSSL Project
sRGB v2 ICC profile, Copyright © International Color Consortium 2016
WOFF2 font decompression, Copyright © 2013-2017 by the WOFF2 Authors
Zlib compression library, Copyright © 1995-2017 Jean-loup Gailly and Mark Adler

The PDFlib Block Plugin contains the following additional third-party component:
wxWidgets Cross-Platform GUI Library, Copyright © 2018 © 1998 Julian Smart, © 2018 wxWidgets



Contents

o Applying the PDFlib License Key 11

1 Introduction 15

- 1.1 What is PDFlib? 15
- 1.2 Roadmap to Documentation and Samples 16
- 1.3 Configuration Requirements 17
- 1.4 What's new in PDFlib/PDFlib+PDI/PPS 10.0? 18
- 1.5 Deprecated and Removed Methods and Options 20
- 1.6 Features in PDFlib 21
- 1.7 Additional Features in PDFlib+PDI and the pCOS Interface 23
- 1.8 Additional Features in PPS and the Block Plugin 24
- 1.9 Availability of Features in different Products 25

2 PDFlib Language Bindings 27

- 2.1 C Binding 27
- 2.2 C++ Binding 30
- 2.3 Java Binding 31
- 2.4 .NET Binding 33
- 2.5 Objective-C Binding 36
- 2.6 Perl Binding 37
- 2.7 PHP Binding 38
- 2.8 Python Binding 40
- 2.9 RPG Binding 41
- 2.10 Ruby Binding 43

3 Creating PDF Documents 47

- 3.1 General PDFlib Programming Aspects 47
 - 3.1.1 Exception Handling 47
 - 3.1.2 Logging 49
 - 3.1.3 The PDFlib Virtual File System (PVF) 49
 - 3.1.4 Resource Configuration and File Search 51
 - 3.1.5 Generating PDF Documents in Memory 55
 - 3.1.6 Maximum Size of PDF Documents and other Limits 56
 - 3.1.7 Web-Optimized (Linearized) PDF 57
 - 3.1.8 Multi-threaded Programming 57
 - 3.1.9 Using PDFlib on EBCDIC-based Platforms 58
- 3.2 Page Descriptions 58
 - 3.2.1 Coordinate Systems 58

- 3.2.2 Page Size **61**
- 3.2.3 Direct Paths **62**
- 3.2.4 Path Objects **62**
- 3.2.5 Templates (Form XObjects) **64**
- 3.3 PDF Password Security 66**
 - 3.3.1 Password Security in PDF **66**
 - 3.3.2 Password-Protecting PDF Documents with PDFlib **66**
- 3.4 XMP Metadata 68**
- 3.5 Associated Files 69**

4 Color and Transparency 73

- 4.1 Device Color Spaces 73**
- 4.2 Color Management with ICC Profiles 74**
 - 4.2.1 Loading ICC Profiles **74**
 - 4.2.2 Using ICC Profiles **75**
 - 4.2.3 Advanced Color Management Controls **77**
- 4.3 Device-Independent CIE L*a*b* Color 78**
- 4.4 Pantone, HKS, and custom Spot Colors 79**
- 4.5 DeviceN Colors 82**
- 4.6 Shadings and Shading Patterns 85**
- 4.7 Tiling Patterns 87**
- 4.8 Transparency Blend Modes 88**
- 4.9 Changing the Color of Objects 90**
 - 4.9.1 Changing the Color with Blend Modes **91**
 - 4.9.2 Changing the Color with Soft Masks **91**
- 4.10 Overprint Control 93**

5 Unicode and Legacy Encodings 95

- 5.1 Important Unicode Concepts 95**
- 5.2 Unicode-capable Language Bindings 96**
 - 5.2.1 Language Bindings with native Unicode Strings **97**
 - 5.2.2 Language Bindings with UTF-8 Support **97**
- 5.3 Non-Unicode-capable Language Bindings 98**
 - 5.3.1 Language Bindings without native Unicode Strings **98**
 - 5.3.2 Unicode Handling and String Types **98**
 - 5.3.3 Single-Byte (8-Bit) Encodings **101**
 - 5.3.4 Chinese, Japanese, and Korean CMaps (deprecated) **103**
- 5.4 Addressing Characters 105**
 - 5.4.1 Escape Sequences **105**
 - 5.4.2 Character References **106**

6 Fonts 109

6.1 Font Formats 109

- 6.1.1 TrueType Fonts 109
- 6.1.2 OpenType Fonts 109
- 6.1.3 OpenType Color Fonts in SVG, COLR and sbix Flavors 110
- 6.1.4 WOFF and WOFF2 Fonts 112
- 6.1.5 PostScript Type 1 Fonts (deprecated) 112
- 6.1.6 SING Fonts (Glyphlets) 112
- 6.1.7 Type 3 Fonts 113

6.2 Unicode Characters and Glyphs 114

- 6.2.1 Glyph IDs 114
- 6.2.2 Unicode Mapping for Glyphs 114

6.3 Loading Fonts 116

- 6.3.1 Selecting an Encoding for Text Fonts 116
- 6.3.2 Selecting an Encoding for symbolic Fonts 117
- 6.3.3 Example: Selecting a Glyph from the Wingdings Symbol Font 119
- 6.3.4 Searching for Fonts 122
- 6.3.5 Host Fonts on Windows and macOS 126
- 6.3.6 Fallback Fonts 127

6.4 Font Embedding and Subsetting 129

- 6.4.1 Embedding 129
- 6.4.2 Subsetting 130

6.5 Querying Font Information 131

- 6.5.1 Font-independent Encoding, Unicode, and Glyph Name Queries 131
- 6.5.2 Font-specific Encoding, Unicode, and Glyph Name Queries 132
- 6.5.3 Querying Codepage Coverage and Fallback Fonts 133

7 Text Output 135

7.1 Text Output Methods 135

7.2 Font Metrics and Text Styles 136

- 7.2.1 Font and Glyph Metrics 136
- 7.2.2 Kerning 136
- 7.2.3 Text Styles 137

7.3 Emoji Sequences 138

7.4 OpenType Layout Features 139

- 7.4.1 Supported OpenType Layout Features 140
- 7.4.2 OpenType Layout Features with Textlines and Textflows 143
- 7.4.3 Script- and Language-specific OpenType Layout Features 143

7.5 Complex Script Output 145

- 7.5.1 Complex Scripts 145
- 7.5.2 Controlling the Shaping Process 146
- 7.5.3 Bidirectional Formatting 147
- 7.5.4 Arabic Text Formatting 149

7.6 Chinese, Japanese, and Korean Text Output 150

- 7.6.1 Using TrueType and OpenType CJK Fonts 150

- 7.6.2 Horizontal and Vertical Writing Mode. **150**
- 7.6.3 EUDC and SING Fonts for Gaiji Characters **151**
- 7.6.4 OpenType Layout Features for CJK Text Output **152**
- 7.7 Unicode Variation Sequences (UVS) 153**
 - 7.7.1 Types of Variation Sequences **153**
 - 7.7.2 Variation Sequences in PDFlib **154**
- 7.8 The Text Processing Pipeline 155**
 - 7.8.1 Normalizing Input Strings to Unicode **155**
 - 7.8.2 Converting Unicode Values to Glyph IDs **156**
 - 7.8.3 Transforming Glyph IDs **157**

8 Importing Images, SVG Graphics and PDF Pages 159

- 8.1 Raster Images 159**
 - 8.1.1 Basic Image Handling **159**
 - 8.1.2 Supported Image File Formats **161**
 - 8.1.3 Clipping Paths **164**
 - 8.1.4 Image Transparency **165**
 - 8.1.5 Colorize Images with Spot or DeviceN Color **168**
 - 8.1.6 Modifying Color Values with a Decode Array **169**
- 8.2 SVG Graphics 170**
 - 8.2.1 Supported SVG Flavors **170**
 - 8.2.2 SVG Processing Considerations **170**
 - 8.2.3 Size of SVG Graphics **172**
 - 8.2.4 Font Selection **173**
 - 8.2.5 Dealing with missing Fonts and missing Glyphs **175**
 - 8.2.6 SVG Color Extension **176**
 - 8.2.7 SVG Contents beyond Vector Graphics and Text **178**
 - 8.2.8 Unsupported SVG Features **179**
- 8.3 Importing PDF Pages with PDI 180**
 - 8.3.1 PDI Features and Applications **181**
 - 8.3.2 Using PDFlib+PDI **181**
 - 8.3.3 Document and Page-related Checks **182**
 - 8.3.4 Importing Annotations, Form Fields and Actions **183**
 - 8.3.5 Additional Aspects of imported PDF Documents **184**
 - 8.3.6 Cloning Document-related Properties **186**
 - 8.3.7 Cloning Page-related Properties **189**
- 8.4 Placing Images, Graphics, and imported PDF Pages 190**
 - 8.4.1 Simple Object Placement **190**
 - 8.4.2 Placing an Object at a Point or Line or in a Box **190**
 - 8.4.3 Orientating an Object **192**
 - 8.4.1 Rotating an Object **193**
 - 8.4.2 Adjusting the Page Size **194**
 - 8.4.1 Querying Information about placed Images and PDF Pages **195**

9 Text and Table Formatting 197

- 9.1 **Placing and Fitting Textlines 197**
 - 9.1.1 Simple Textline Placement 197
 - 9.1.1 Positioning Text in a Box 198
 - 9.1.2 Fitting Text into a Box 199
 - 9.1.3 Aligning Text at a Character 201
 - 9.1.4 Placing a Stamp 202
 - 9.1.5 Using Leaders 202
 - 9.1.6 Text on a Path 203
 - 9.1.7 Shadowed Text 204
- 9.2 **Multi-Line Textflows 205**
 - 9.2.1 Placing Textflows in the Fitbox 206
 - 9.2.2 Paragraph Formatting Options 208
 - 9.2.1 Inline Option Lists and Macros 209
 - 9.2.1 Tab Stops 211
 - 9.2.1 Numbered Lists and Paragraph Spacing 212
 - 9.2.1 Control Characters and Character Mapping 214
 - 9.2.2 Hyphenation 216
 - 9.2.3 Widow and Orphan Lines 216
 - 9.2.4 Controlling the standard Linebreak Algorithm 217
 - 9.2.5 Advanced script-specific Line Breaking 221
 - 9.2.6 Wrapping Text around Paths and Images 221
- 9.3 **Table Formatting 225**
 - 9.3.1 Placing a Simple Table 227
 - 9.3.2 Contents of a Table Cell 230
 - 9.3.3 Table and Column Widths 232
 - 9.3.4 Mixed Table Contents 233
 - 9.3.5 Table Instances 235
 - 9.3.6 Table Formatting Algorithm 238
- 9.4 **Matchboxes 241**
 - 9.4.1 Decorating a Textline 241
 - 9.4.2 Using Matchboxes in a Textflow 242
 - 9.4.3 Matchboxes and Images 243

10 Interactive Features 245

- 10.1 **Links, Bookmarks, and Annotations 245**
- 10.2 **Form Fields and JavaScript 247**
- 10.3 **Multimedia with Screen Annotations and Rendition Actions 249**
- 10.4 **Geospatial PDF 252**
 - 10.4.1 Using georeferenced PDF in Acrobat 252
 - 10.4.2 Geographic and projected Coordinate Systems 253
 - 10.4.3 Coordinate System Examples 253
 - 10.4.4 Georeferenced PDF Restrictions in Acrobat 254

11 Tagged PDF and PDF/UA 257

- 11.1 Tagged PDF Basics 257**
 - 11.1.1 Creating the Logical Structure Tree (Structure Hierarchy) **257**
 - 11.1.2 Standard Structure Types **261**
 - 11.1.3 PDF 1.7 Structure Element Nesting Rules **266**
 - 11.1.4 PDF 2.0 Structure Element Nesting Rules **268**
 - 11.1.5 Custom Structure Types, Role Maps and custom Tagsets **271**
 - 11.1.6 Artifacts **272**
 - 11.1.7 Text Handling **274**
 - 11.1.8 Alternate Description, Replacement Text and Abbreviation Expansion **275**
 - 11.1.9 Print Stream Order and Logical Reading Order **276**
 - 11.1.1 Tagged PDF Problems in Adobe Acrobat **278**
- 11.2 Advanced Tagged PDF Creation Topics 278**
 - 11.2.1 Automatic Table Tagging **278**
 - 11.2.1 Tagging Interactive Elements **281**
 - 11.2.2 PDF 2.0 Features for Tagging Interactive Elements **286**
 - 11.2.3 Tagging Lists **286**
 - 11.2.4 Creating Contents out of Order **288**
- 11.3 Importing Tagged PDF Pages with PDI 290**
- 11.4 PDF/UA for Universal Accessibility 293**
 - 11.4.1 The PDF/UA-1 Standard **293**
 - 11.4.2 Tagging Requirements **296**
 - 11.4.3 Additional Requirements for specific Content Types **298**
 - 11.4.4 Importing Pages from PDF/UA Documents with PDI **298**

12 PDF Versions and Standards 301

- 12.1 PDF Versions 301**
 - 12.1.1 PDF Version Overview **301**
 - 12.1.2 PDF 1.x Features **301**
 - 12.1.3 PDF 2.0 Features **303**
 - 12.1.4 Deprecated PDF Features **304**
- 12.2 PDF/A for Archiving 305**
 - 12.2.1 The PDF/A Standards **305**
 - 12.2.2 General Requirements **306**
 - 12.2.3 Color and Image Requirements **307**
 - 12.2.4 Requirements for Interactive Features **310**
 - 12.2.5 Additional PDF/A Requirements for Level U Conformance **311**
 - 12.2.6 Additional PDF/A Requirements for Level A Conformance **311**
 - 12.2.7 Importing PDF/A Documents with PDI **312**
 - 12.2.8 XMP Metadata for PDF/A **314**
- 12.3 PDF/X for Print Production 316**
 - 12.3.1 The PDF/X Family of Standards **316**
 - 12.3.2 General Requirements **317**
 - 12.3.3 Output Intent and Color Requirements **318**
 - 12.3.4 Image and Transparency Requirements **321**
 - 12.3.5 Requirements for interactive Features **322**

- 12.3.6 Importing PDF/X Documents with PDI 322
- 12.4 PDF/VT for Variable and Transactional Printing 324
 - 12.4.1 The PDF/VT Standard 324
 - 12.4.2 PDF/VT Concepts 324
 - 12.4.3 Summary of Rules for generating PDF/VT 326
 - 12.4.4 Document Part Hierarchy and Document Part Metadata (DPM) 327
 - 12.4.5 Scope Hints for recurring Graphical Content 328
 - 12.4.6 Encapsulated XObjects 329
 - 12.4.7 Importing PDF/X and PDF/VT Documents with PDI 330

13 PPS and the PDFlib Block Plugin 333

- 13.1 Installing the PDFlib Block Plugin 333
- 13.2 Overview of the Block Concept 335
 - 13.2.1 Separation of Document Design and Program Code 335
 - 13.2.2 Block Properties 335
 - 13.2.3 Why not use PDF Form Fields? 336
- 13.3 Editing Blocks with the Block Plugin 337
 - 13.3.1 Creating Blocks 337
 - 13.3.2 Editing Block Properties 341
 - 13.3.3 Block Filtering 343
 - 13.3.4 Copying Blocks between Pages and Documents 344
 - 13.3.5 Customizing the Block Plugin User Interface with XML 345
- 13.4 Converting PDF Form Fields to PDFlib Blocks 346
- 13.5 Previewing Blocks in Acrobat 349
- 13.6 Filling Blocks with PPS 353
- 13.7 Block Properties 357
 - 13.7.1 Administrative Properties 358
 - 13.7.2 Rectangle Properties 359
 - 13.7.3 Appearance Properties 360
 - 13.7.4 Text Preparation Properties 362
 - 13.7.5 Text Formatting Properties 363
 - 13.7.6 Object Fitting Properties 366
 - 13.7.7 Tagged PDF Properties 368
 - 13.7.8 Properties for default Contents 369
 - 13.7.9 Custom Properties 369
- 13.8 Querying Block Names and Properties with pCOS 369
- 13.9 Creating and Importing Blocks programmatically 371
 - 13.9.1 Creating PDFlib Blocks with POCA 371
 - 13.9.2 Importing PDFlib Blocks 373
- 13.10 PDFlib Block Specification 373

A Revision History 377

Index 379

o Applying the PDFlib License Key

Restrictions of the evaluation version. All binary versions of PDFlib, PDFlib+PDI, and PPS supplied by PDFlib GmbH can be used as fully functional evaluation versions regardless of whether or not you obtained a commercial license. However, unlicensed versions display a *www.pdflib.com* demo stamp across all generated pages, and the integrated pCOS interface is limited to small documents (up to 10 pages and 1 MB file size). Unlicensed binaries must not be used for production purposes, but only for evaluating the product. Using any PDFlib GmbH product for production purposes requires a valid license.

Companies which are interested in PDFlib licensing and wish to get rid of the evaluation restrictions during the evaluation phase or for prototype demos can submit their company and project details with a brief explanation to *sales@pdflib.com*, and apply for a temporary license key (we reserve the right to refuse evaluation key requests, e.g. for anonymous requests).

PDFlib, PDFlib+PDI, and PDFlib Personalization Server (PPS) are different products which require different license keys although they are delivered in a single package. PDFlib+PDI license keys will also be valid for PDFlib, but not vice versa, and PPS license keys will be valid for PDFlib+PDI and PDFlib. All license keys are platform-dependent, and can only be used on the platform for which they have been purchased.

Once you purchased a license key you must apply it in order to get rid of the demo stamp. Several methods are supported for setting the license key; they are detailed below.

Cookbook A full code sample can be found in the *Cookbook* topic `general/license_key`.

Applying a license key with an API call at runtime. Add a line to your script or program which sets the license key at runtime. The *license* option must be set immediately after instantiating the PDFlib object (in C: after *PDF_new()*). The exact syntax depends on your programming language:

- ▶ In C++, Java, .NET/C#, Python and Ruby:

```
p.set_option("license=...your license key...")
```

- ▶ In C:

```
PDF_set_option(p, "license=...your license key...")
```

- ▶ In Objective-C:

```
[pdflib set_option: @"license=...your license key..."];
```

- ▶ In Perl and PHP:

```
$p->set_option("license=...your license key...")
```

- ▶ In RPG:

```
c                                callp      PDF_set_option(p:%ucs2('license=...your  
license key...')
```

Working with a license file. As an alternative to supplying the license key with a runtime call, you can enter the license key in a text file according to the following format

(you can use the license file template *licensekeys.txt* which is contained in all PDFlib distributions). Lines beginning with a '#' character contain comments and will be ignored; the second line contains version information for the license file itself:

```
# Licensing information for PDFlib GmbH products
PDFlib license file 1.0
PDFlib      10.0.1      ...your license key...
```

The license file may contain license keys for multiple PDFlib GmbH products on separate lines. It may also contain license keys for multiple platforms so that the same license file can be shared among platforms. License files can be configured in the following ways:

- ▶ A file called *licensekeys.txt* will be searched in all default locations (see »Default file search paths«, page 12).
- ▶ You can specify the *licensefile* option with the *set_option()* API method:

```
p.set_option("licensefile={/path/to/licensekeys.txt}");
```

- ▶ You can set an environment (shell) variable which points to a license file. On Windows use the system control panel and choose *System, Advanced, Environment Variables.*; on Unix apply a command similar to the following:

```
export PDFLIBLICENSEFILE=/path/to/licensekeys.txt
```

- ▶ On IBM System i the license file must be encoded in ASCII (see *asciifile* option). The license file can be specified as follows (this command can be specified in the startup program *QSTRUP* and works for all PDFlib GmbH products):

```
ADDENVVAR ENVVAR(PDFLIBLICENSEFILE) VALUE('/PDFlib/10.0/licensefile.txt') LEVEL(*SYS)
```

License keys in the registry. On Windows you can also enter the name of the license file in the following registry value:

```
HKLM\SOFTWARE\PDFlib\PDFLIBLICENSEFILE
```

As another alternative you can enter the license key directly in one of the following registry values:

```
HKLM\SOFTWARE\PDFlib\PDFlib10\license
HKLM\SOFTWARE\PDFlib\PDFlib10\10.0.1\license
```

Note Be careful when manually accessing the registry on 64-bit Windows systems: as usual, 64-bit PDFlib binaries work with the 64-bit view of the Windows registry, while 32-bit PDFlib binaries running on a 64-bit system work with the 32-bit view of the registry. If you add registry keys for a 32-bit product manually, make sure to use the 32-bit version of the *regedit* tool. It can be invoked as follows from the Start dialog:

```
%systemroot%\syswow64\regedit
```

Default file search paths. On Unix, Linux, macOS and IBM System i some directories will be searched for files by default even without specifying any path and directory names. Before searching and reading the UPR file (which may contain additional search paths), the following directories will be searched:

```
<rootpath>/PDFlib/PDFlib/10.0/resource/cmap
<rootpath>/PDFlib/PDFlib/10.0/resource/font
<rootpath>/PDFlib/PDFlib/10.0/resource/icc
<rootpath>/PDFlib/PDFlib/10.0
<rootpath>/PDFlib/PDFlib
<rootpath>/PDFlib
```

On Unix, Linux, and macOS *<rootpath>* is first replaced with */usr/local* and then with the HOME directory. On IBM System i *<rootpath>* is empty.

Default file names for license and resource files. By default, the following file names will be searched for in the default search path directories:

licensekeys.txt	(license file)
pdflib.upr	(resource file)

This feature can be used to work with a license file without setting any environment variable or runtime option.

Updates and Upgrades. If you purchased an update (change from an older version of a product to a newer version of the same product) or an upgrade (change from PDFlib to PDFlib+PDI or PPS, or from PDFlib+PDI to PPS), or received a new license key as part of your support contract, you must apply the new license key that you received for your update or upgrade. The old license key for the previous product must no longer be used.

Evaluating features which are not yet licensed. You can fully evaluate all features by using the software without any license key applied. However, once you applied a valid license key for a particular product using features of a higher category will no longer be available. For example, if you installed a valid PDFlib license key the PDI functionality will no longer be available for testing. Similarly, after installing a PDFlib+PDI license key the personalization features (block functions) will no longer be available.

When a license key for a product has already been installed, you can replace it with the dummy license string »0« (digit zero) to enable functionality of a higher product class for evaluation. This will enable the previously disabled functions, and re-activate the demo stamp across all pages. License key »0« does not work in PDFlib ME.

Licensing options. Different licensing options are available for PDFlib use on one or more servers and for redistributing PDFlib with your own products. We also offer support and source code contracts. Licensing details and the PDFlib purchase order form can be found in the PDFlib distribution. Please contact us if you are interested in obtaining a PDFlib license or have any questions.

1 Introduction

1.1 What is PDFlib?

PDFlib is a well-established developer component for generating and manipulating PDF documents. PDFlib's main targets are dynamic PDF creation on a Web server and PDF creation in any kind of server, desktop or mobile application. You can use PDFlib to dynamically create PDF documents from database contents, similar to dynamic Web pages. Since 1997 PDFlib has proven itself in a wide range of use cases and development environments. Application programmers use PDFlib to create PDF output with text, images, graphics and interactive elements.

PDFlib is in active use by tens of thousands of users worldwide. The large-scale deployment helps us to improve and fine-tune the API, the internal implementation and the generated PDF output. Conformance to the relevant ISO standards for PDF, PDF/A, PDF/X, PDF/UA and PDF/VT as well as compatibility with common PDF viewers have high priority. PDFlib supports all features provided by the PDF format.

The PDFlib product family is available in three different flavors: PDFlib, PDFlib+PDI (PDF Import), and the PDFlib Personalization Server (PPS) with the PDFlib Block Plugin for Adobe Acrobat.

PDFlib. PDFlib offers methods for generating PDF documents with text, graphics, images, and interactive elements such as annotations or bookmarks. The PDFlib API supports PDF creation on different levels which can freely be mixed.

Low-level APIs allow you to place individual items on the page, using formatting information provided by the application. An application which already determined the layout (e.g. for screen display) can directly place text, images and other objects on the PDF page.

High-level APIs support powerful formatting features:

- ▶ Single-line text output with various formatting controls.
- ▶ Textflow formatter for placing arbitrary amounts of text in one or more columns or pages. The text can be formatted in non-rectangular shapes, flow around images and contain interactive elements and structure information for Tagged PDF.
- ▶ The flexible table formatter supports all types of contents in table cells like text, images, graphics and annotations or form fields. Large tables can be placed on multiple pages. Details of table formatting are controlled by a variety of options.
- ▶ Image and SVG placement according to various fitting and formatting methods.

PDFlib+PDI (PDF Import). PDFlib+PDI includes all PDFlib functions plus the PDF Import Library PDI which incorporates existing PDF pages and interactive elements in the generated PDF output. Use PDFlib+PDI for all PDFlib tasks plus the following:

- ▶ impose multiple PDF pages on a single sheet
- ▶ add text, such as headers, footers, stamps, or page numbers to existing PDF pages
- ▶ place images, e.g. company logo, on existing pages
- ▶ rearrange, assemble or amend PDF pages while maintaining PDF/A, PDF/X or PDF/UA conformance

PDFlib Personalization Server (PPS) and PDFlib Block Plugin. The PDFlib Personalization Server (PPS) includes PDFlib+PDI plus additional methods for variable data processing using PDFlib Blocks. PPS makes applications independent from layout changes.

The designer creates a layout PDF and marks areas as placeholders for variable text, images and graphics using the PDFlib Block Plugin for Acrobat. Each Block contains a variety of Block properties such as font size, color, image scaling options etc. The PDFlib Block Plugin offers a Preview feature which shows the results of filling Blocks according to the specified properties.

The developer writes code to fill PDFlib Blocks with text, images, vector graphics or PDF pages. He doesn't need to know the formatting or position of a Block. Use PPS for all PDFlib+PDI tasks plus mail merge with text and images, template filling for transactional and statement processing, personalizing promotional material, generating individual parts catalogs from a database or producing customized documentation for multiple products.

How can I use PDFlib? PDFlib is available on a variety of platforms including Unix, Windows, macOS, and EBCDIC-based systems such as IBM System i and IBM Z. PDFlib is written in C and C++, but it can also be accessed from several other languages and programming environments which are called language bindings. These bindings cover all current Web and stand-alone development environments. The Application Programming Interface (API) is easy to learn, and is identical for all bindings. Currently the following bindings are supported:

- ▶ C and C++
- ▶ Java
- ▶ .NET
- ▶ Objective-C (macOS and iOS) and Swift
- ▶ Perl
- ▶ PHP
- ▶ Python
- ▶ RPG (IBM System i)
- ▶ Ruby

1.2 Roadmap to Documentation and Samples

We provide the material listed below to assist you in using PDFlib products successfully.

Mini samples for all language bindings. The *hello* and *image* mini samples are available in all packages and for all language bindings. They provide minimal sample code for text and image output. The mini samples are mainly useful for testing your PDFlib installation and for getting a very quick overview of writing PDFlib application code.

Starter samples for all language bindings. The *starter samples* are contained in all packages and are available for all language bindings. They provide a useful starting point for important topics and cover simple text and image output, Textflow and table formatting, PDF/A, PDF/X and PDF/UA creation as well as many other topics. The starter samples demonstrate basic techniques for achieving a particular goal with PDFlib products. It is strongly recommended to take a look at the starter samples.

PDFlib Tutorial. The *PDFlib Tutorial* (this manual), which is included in all packages, explains important programming concepts in more detail, including small code fragments. If you start extending your code beyond the starter samples you should read up on relevant topics in the PDFlib Tutorial.

Note Most examples in this PDFlib Tutorial are provided in the Java language. Although syntax details vary with each language, the basic concepts of PDFlib programming are the same for all language bindings.

PDFlib API Reference. The *PDFlib API Reference*, which is contained in all packages, contains a concise description of all functions and options which together comprise the PDFlib application programming interface (API). The PDFlib API Reference is the definitive source for looking up supported options, input conditions, and other programming rules which must be obeyed.. Make sure to always use the PDFlib API Reference when developing with PDFlib.

pCOS Path Reference. The pCOS interface can be used to query a variety of properties from PDF documents. pCOS is included in PDFlib+PDI and PPS. The pCOS Path Reference contains a description of the path syntax used to address individual objects within a PDF document in order to retrieve the corresponding values.

PDFlib Cookbook. The *PDFlib Cookbook* contains hundreds of PDFlib coding fragments for specific tasks. The Cookbook topics are available for Java and PHP, but can easily be adjusted to other programming languages since the PDFlib API is identical for all supported language bindings. The PDFlib Cookbook is available at the following URL:

www.pdflib.com/pdflib-cookbook/

pCOS Cookbook. The *pCOS Cookbook* is a collection of Java code fragments for the pCOS interface which is contained in PDFlib+PDI and PPS. It is available at the following URL:

www.pdflib.com/pcos-cookbook/

TET Cookbook. PDFlib TET (Text and Image Extraction Toolkit) is a separate product for extracting text and images from PDF documents. It can be combined with PDFlib+PDI to process PDF documents based on their contents. The *TET Cookbook* is a collection of Java and XSLT code fragments for TET. It contains a group of samples which demonstrate the combination of TET and PDFlib+PDI, e.g. add Web links or bookmarks based on the text on the page, highlight search terms, split documents based on text, create a table of contents, etc. The TET Cookbook is available at the following URL:

www.pdflib.com/tet-cookbook/

1.3 Configuration Requirements

Unlike earlier versions PDFlib 10 requires some external files for specific operations. The required files are included in the PDFlib package. It is up to the user to copy these files to suitable directories and to configure access to these directories. Several methods including *searchpath* are available for configuring access to external files; see Section

3.1.4, »Resource Configuration and File Search«, page 51, for details. External files are required for the operations listed in Table 1.1.

Table 1.1 Configuration requirements

<i>files</i>	<i>required for</i>
resource/codepage/*.cpg	<i>PDF_load_font()</i> and implicit font loading with one of the legacy 8-bit encodings iso8859-2...iso8859-16, cp1250...cp1251 and cp1253...cp1258 (see Section 5.3.3, »Single-Byte (8-Bit) Encodings«, page 101)
resource/font/*.ttf	<i>PDF_create_annotation()</i> with type=Caret, FreeText, Line or Stamp if neither the font nor template option is supplied. <i>PDF_fit_pdi_page()</i> with useannots (enabled by default) or usefields if a page containing form fields or one of the annotation types above is imported and the annotation or field doesn't have an appearance stream
resource/cmap/*	<i>PDF_load_font()</i> and implicit font loading with one of the legacy CJK encodings (CMaps). Note that these encodings are deprecated (see Section 5.3.4, »Chinese, Japanese, and Korean CMaps (deprecated)«, page 103)

1.4 What's new in PDFlib/PDFlib+PDI/PPS 10.0?

The following provides an overview of the most important new or improved features in PDFlib/PDFlib+PDI/PPS 10.0 and Block Plugin 6. There are many more new features; see Table 1.2 and the PDFlib API Reference for details.

Color fonts and emoji. OpenType color fonts allow glyph descriptions with color and transparency. They are especially popular for emoji fonts. PDFlib supports color fonts based on SVG or Microsoft COLR glyphs along with OpenType features, variation selectors and emoji variation sequences, e.g. for changing skin color or gender of a particular emoji.

Import form fields and annotations with PDFlib+PDI. PDFlib+PDI traditionally imported only the page contents, but ignored interactive elements. Links and other annotation types, form fields, actions and JavaScript are imported along with the page contents to ensure that interactive documents retain their functionality. Interactive elements remain functional even if the imported pages are placed in different order or if multiple imported pages are placed on the same output page. Annotation rectangles are transformed appropriately if the page containing the annotation has been scaled or rotated. Annotation and form field import are compatible with PDF/A, PDF/UA and PDF/X.

Annotations. Creation of all types of annotations (comments) has been overhauled, extended and simplified. Annotation appearance streams are created automatically; annotations are supported in PDF/A mode as far as allowed by the standard.

Form filling. Imported form fields can be filled, e.g. to populate text fields. The contents and visual appearance (color, border etc.) of imported form fields can be modified or kept in the original state.

Multimedia. Video and sound files can be used in PDF with Screen annotations and Rendition actions. These are more powerful than Sound and Movie annotations and

don't require the deprecated Flash technology which is no longer supported in PDFlib 10 nor Acrobat DC.

PDF 2.0 support. PDFlib supports new features in PDF 2.0 according to ISO 32000-2:2017 and the dated revision ISO 32000-2:2020. This includes new structure element types and nesting rules for Tagged PDF, graphical features, encryption, interactive elements, document and page-level output intents and many other areas. PDFlib+PDI and PPS process PDF 2.0 documents as well. Features marked as deprecated in PDF 2.0 are also deprecated or removed in PDFlib.

Color management. The industry-standard sRGB ICC profile is automatically applied to RGB images which don't include an embedded ICC profile. This facilitates faithful color reproduction and use of RGB images in PDF/A or PDF/X documents.

Faithful color reproduction of imported PDF pages and SVG graphics with transparency has been simplified by automatic creation of transparency groups.

Optimized PDF output. ICC profiles and other objects in PDF are cached and duplicates are detected. In combination with optimized content streams with fewer operators this results in significant PDF file size reductions. Large PDF documents beyond 10 GB can be created in combination with optimization and linearization.

SVG import. Several details of SVG processing have been improved including certain CSS directives. Remote resources, e.g. fonts and images, can be fetched over the network.

Font support and text handling. Enhancements in addition to color font support:

- ▶ Fonts in WOFF2 format can be used.
- ▶ Improved support for OpenType layout features.
- ▶ Improved Type 3 font creation.

Convenience features. Many existing API methods have been enhanced with convenience options and simplified procedures for common tasks. For example, graphics state properties can be supplied directly via options to many API methods without the need for creating and applying a gstate object. Path objects can be drawn inside a matchbox which has been created earlier, e.g. via Textflow or Table formatting.

The default PDF output compatibility is PDF 1.7 Extension Level 8, the file format of Acrobat X/XI/DC.

Updated Pantone® color libraries. The integrated Pantone spot color database has been extended to cover the latest additions to the *Solid Coated* and *Solid Uncoated* color libraries. Spot colors which are no longer supported by Pantone, Inc. have been removed, e.g. the hexachrome color library.

pCOS interface. The pCOS interface (which is included in PDFlib+PDI and PPS) supports additional pseudo objects for querying PDF details, e.g. new PDF standards, form fields, signatures and ICC profiles.

Tagged PDF and PDF/UA. Automatic Table tagging tags tables which span multiple pages correctly with a single Table element. Links spanning multiple lines are created as

a single Link annotation with multiple rectangles. Structure element tags can be supplied inline in Textflow along with formatting options.

PDF 2.0 introduces a strict model for the relationship of structure elements. PDFlib supports new PDF 2.0 structure types and attributes and enforces the new structure element nesting rules. Tag namespaces (tagsets) and structure destinations are supported.

Untagged documents can be imported into PDF/UA as Artifact which simplifies handling of existing decorative content which is not available as PDF/UA.

PDFlib Personalization Server (PPS) and Block Plugin. The PDFlib Block Plugin for Acrobat received various usability improvements including a filter for temporarily reducing the number of visible Blocks for screen and Preview. Blocks can easily be moved by small or large amounts. Additional Block properties support all graphics state options including overprint settings. Structure elements (tags) are available as Block property and Preview supports PDF/UA cloning.

Networking. Remote resources can be fetched from the network with a new PDFlib API method, i.e. independently of the means provided by the environment. SVG processing automatically retrieves remote resources such as fonts or images.

Language bindings. All language bindings have been updated to the latest versions. The Perl, PHP and Ruby bindings are UTF-8-enabled by default. The C++ binding supports the Unicode string types *u8string/u16string/u32string* introduced with C++11 and C++20.

Miscellaneous improvements. New features in various areas:

- ▶ Type 3 fonts may contain a mixture of colorized and non-colorized glyphs.
- ▶ Additional appearance options for path objects.

Modernized code base and improved code security. The code has been reviewed and improved for robustness; unnecessary code has been removed, e.g. workarounds for outdated versions of Acrobat and RIPS. Removal of deprecated features also resulted in smaller and cleaner code. Third-party libraries used in PDFlib have been updated to the latest versions to take advantage of security improvements and vulnerability protection.

1.5 Deprecated and Removed Methods and Options

While a number of PDFlib API methods have been declared as deprecated in PDFlib 9 and earlier versions, API methods have never been removed so far, but were still offered as migration aid. PDFlib 10 removes all API methods which have been declared as deprecated in PDFlib 9 or earlier versions. In most cases simple one-to-one replacements are available.

Similarly, deprecated options have been removed. In most cases the underlying functionality is still supported and simple replacement options are available.

Some features are no longer available at all, either because there is no demand in the marketplace (e.g. Reference XObjects) or because they have been declared as deprecated in PDF 2.0, e.g. OPI and Flash.

PDFlib 9 applications which don't use any deprecated features can be used unmodified with PDFlib 10. If your application uses deprecated features please take a look at the

PDFlib Migration Guide. This document contains detailed lists of deprecated and removed API methods, options and features as well as notes regarding the required changes in application code.

The PDFlib Migration Guide also explains how to identify deprecated features with PDFlib 9 before migrating to PDFlib 10.

1.6 Features in PDFlib

Table 1.2 lists features for generating PDF. New and improved features are marked. Note that improved PDF output (e.g. smaller files) and programming convenience features are not listed individually.

Table 1.2 Feature list for PDFlib

topic	features
PDF flavors	PDF 1.4 – PDF 1.7 extension level 8 and PDF 2.0 ¹
	Linearized (web-optimized) PDF for byteserving over the Web
	High-volume output and PDF file size beyond 10 GB
ISO standards for PDF	ISO 32000-1: standardized version of PDF 1.7
	ISO 32000-2 including dated revision ISO 32000-2:2020 ¹ : PDF 2.0
	ISO 15930: PDF/X-3/4/5 for the graphic arts industry
	ISO 19005-1/2/3: PDF/A-1/2/3 for archiving
	ISO 16612-2: PDF/VT-1 for variable and transactional printing
ISO 14289-1: PDF/UA-1 for universal accessibility	
Fonts	TrueType fonts (TTF and TTC)
	OpenType fonts with PostScript or TrueType outlines (TTF, OTF, OTC)
	OpenType color fonts in the formats SVG, COLR (Microsoft) and sbix (Apple)
	WOFF and WOFF2 ¹ fonts (Web Open Font Format)
	User-defined (Type 3) fonts for bitmap, logo and color fonts
	Support for OpenType layout features ¹ for Western and CJK text, e.g. ligatures, small caps, old-style numerals, swash characters, simplified/traditional forms, vertical alternates
	Access fonts which are installed on Windows or macOS
	Font embedding and subsetting
EUDC and SING fonts (glyphlets) for CJK Gaiji characters	
Fallback fonts: substitute missing glyphs with glyphs from another font	
Text output	Text output in different fonts; underlined, overlined, and strikethrough text
	Glyphs in a font can be addressed by numerical value, Unicode value, or glyph name
	Kerning for improved character spacing
	Variation selectors and emoji variation sequences ¹
	Artificial bold, italic, and shadow text
Text on a path	
Flexible stamping method	
Textflow Formatting	Format text into one or more rectangular or arbitrarily shaped areas with hyphenation (user-supplied hyphenation points required), font and color changes, justification methods, tabs, leaders, control commands, tags ¹
	Advanced line-breaking with language-specific processing, e.g. for Thai
	Flexible image placement and formatting

Table 1.2 Feature list for PDFlib

topic	features
	Wrap text around images or image clipping paths
Table formatting	Table formatter places rows and columns, and automatically calculates their sizes according to a variety of user preferences. Tables can be split across multiple pages. Table cells can hold text, images, SVG graphics, PDF pages, path objects, annotations and form fields Table cells can be formatted with ruling and shading options Matchbox concept for referencing the coordinates of placed images or other objects
Interactive elements	Form fields ¹ with all field options and JavaScript PDF annotations ¹ (comments) such as PDF links, launch links (other document types), Web links PDF 2.0 annotation attributes, e.g. opacity ¹ Non-rectangular annotations and Link annotations with multiple rectangles ¹ Create actions for bookmarks, annotations, page open/close and other events Create bookmarks with a variety of options and controls Page transition effects, such as shades and mosaic Named destinations for links, bookmarks, and document open action Create page labels (symbolic names for pages)
Multimedia	Embed 3D animations in PDF and control it with JavaScript Sound and video can be embedded or referenced from an external file or Web server with Screen annotations and Rendition actions ¹
Internationalization	Support for Unicode 14 including supplementary characters outside the Basic Multilingual Plane (BMP) and Unicode normalization forms CJK fonts for Chinese, Japanese, and Korean text Ideographic variation sequences (IVS) for CJK variant glyphs Vertical writing mode for Chinese, Japanese and Korean text Character shaping for complex scripts, e.g. Arabic, Thai, Devanagari Bidirectional text formatting for right-to-left scripts, e.g. Arabic and Hebrew
SVG vector graphics	Import vector graphics in SVG format; ICC profiles, CMYK and spot colors in SVG, CSS ¹ , network resources ¹
Images	Load BMP, PNG, TIFF, JBIG2, JPEG, and JPEG 2000 raster images Query image information (pixel size, resolution, ICC profile, clipping path, etc.) Use clipping path in TIFF and JPEG images Use alpha channel (transparency) in TIFF and PNG images Image masks (transparent images with a color applied), colorize images with spot or DeviceN color
Color	Grayscale, RGB (numerical, hexadecimal, HTML color names), CMYK, CIE L*a*b* color Integrated PANTONE® and HKS® spot colors ¹ DeviceN (n-colorant) color space based on process or spot colors User-defined spot colors Color shadings (smooth shadings) between process colors or spot colors; pattern fills and strokes
Color management	ICC-based color with ICC profiles ¹ Rendering intent for text, graphics, and raster images Embedded or externally referenced ICC profiles as output intent for PDF 2.0, PDF/A and PDF/X Multicolor (xCLR) externally referenced output intent ICC profiles for PDF/X-5n Output intent ICC profiles on document or page level for PDF 2.0

Table 1.2 Feature list for PDFlib

topic	features
Archiving	PDF/A-1a/b, PDF/A-2a/b/u and PDF/A-3a/b/u including XMP extension schemas
Graphic arts	PDF/X-3, PDF/X-4, PDF/X-4p, PDF/X-5n Settings for overprint, text knockout, black point compensation ¹ , etc.
Accessibility	Tagged PDF for accessibility Tagging of interactive elements, e.g. annotations and form fields ¹ Automatic table ¹ and Artifact tagging PDF/UA-1 for universal accessibility Tagging in Textflow ¹ PDF 2.0 tag namespaces (tagsets) ¹ PDF 2.0 structure types, attributes and tag nesting rules ¹ Structure destinations, i.e. links point to a specific structure element ¹
Variable Document Printing (VDP)	PDF/VT-1 for variable and transactional printing
Vector graphics	Common vector graphics primitives ¹ : lines, curves, arcs, ellipses, rectangles, etc. Smooth shadings (color blends) between multiple process or spot colors, pattern fills and strokes Transparency (opacity) and blend modes Reusable path objects ¹ and clipping paths imported from images
Layers	Optional page content which can selectively be displayed Annotations and form fields can be placed on layers
Security	Encrypt PDF document or attachments Unicode passwords Document permission settings, e.g. printing or copying not allowed
Georeferenced PDF	PDF with geospatial reference information
Metadata	Document information: common fields (Title, Subject, Author, Keywords) and user-defined fields Synchronize XMP metadata with document info fields User-supplied XMP metadata Process XMP image metadata in TIFF, JPEG, JPEG 2000 and SVG graphics
Programming	Language bindings ¹ for C, C++, Java, .NET, Objective-C, Perl, PHP, Python, RPG, Ruby Helper method for downloading content from the Internet ¹ Virtual file system for supplying data in memory, e.g., images from a database Generate PDF documents on disk file or directly in memory

1. New or considerably improved in PDFlib 10

1.7 Additional Features in PDFlib+PDI and the pCOS Interface

Table 1.3 lists features in PDFlib+PDI and PPS in addition to the basic PDF generation features in Table 1.2.

Table 1.3 Additional features in PDFlib+PDI

topic	features
PDF input (PDI)	Import pages from existing PDF documents Import all PDF versions up to PDF 2.0

Table 1.3 Additional features in PDFlib+PDI

topic	features
	Import annotations and form fields including actions ¹ ; annotation rectangles are transformed, e.g. scaled or rotated, to match the placed page; link annotations remain functional even if pages are arranged in different order
	Import JavaScript and other actions on document or page level ¹
	Fill imported form fields ¹
	Import encrypted documents
	Delete redundant objects (e.g. identical fonts) across multiple imported PDF documents
	Repair malformed input PDF documents
	Copy PDF/A or PDF/X output intent from imported PDF documents
	Tagged PDF content aggregation: import pages from Tagged PDF documents including structure elements
	Import layer definitions (optional content)
	Improved performance for importing content from very large documents ¹
pCOS interface	pCOS interface for querying details about PDF documents and pages
	Additional pCOS pseudo objects for querying PDF details, e.g. PDF standards, form fields, signatures and ICC profiles ¹

1. New in PDFlib+PDI 10

1.8 Additional Features in PPS and the Block Plugin

Table 1.4 lists features which are only available in the PDFlib Personalization Server (PPS) in addition to the basic PDF generation features in Table 1.2 and the PDF import features in Table 1.3.

Table 1.4 Additional features in the PDFlib Personalization Server (PPS)

topic	features
Variable Document Printing (VDP)	PDF personalization with PDFlib Blocks for text, image, PDF data or SVG vector graphics
	Create PDFlib Blocks programmatically with PPS
	Copy PDFlib Blocks from imported documents
PDFlib Block Plugin	PDFlib Block Plugin for creating PDFlib Blocks interactively in Acrobat on Windows and macOS
	Preview PPS Block filling in Acrobat
	Filter Blocks for Display and Preview based on Block name, type, size or position ¹
	Copy Blocks to Preview file
	Snap-to-grid for interactively creating or editing Blocks in Acrobat
	Move Blocks by small or large increments for simple yet precise positioning ¹
	All graphics state properties, e.g. overprint settings, available as Block properties ¹
	Clone PDF/A, PDF/UA or PDF/X status of the Block container
	Convert PDF form fields to PDFlib Blocks for automated filling
	Textflow Blocks can be linked so that one Block holds the overflow text of a previous Block
	PANTONE® and HKS® spot color names integrated in the Block plugin ¹
	Tagging for Block contents ¹

1. New or improved in Block Plugin 6

1.9 Availability of Features in different Products

Table 1.5 details the availability of features in different products in the PDFlib family.

Table 1.5 Availability of features in different products

feature	API methods and options	PDFlib	PDFlib+PDI	PPS
Basic PDF generation	all except those listed below	X	X	X
Linearized (Web-optimized) PDF	linearize option in <code>PDF_begin_document()</code>	X ¹	X	X
Optimize PDF (for inefficient client code and non-optimized imported PDF documents)	optimize option in <code>PDF_end_document()</code>	X ¹	X	X
Parsing attachments for PDF/A conformance	<code>PDF_load_asset()</code> with type=Attachment in PDF/A-2 mode	X ¹	X	X
Parse PDF documents for Portfolio creation	password option in <code>PDF_add_portfolio_file()</code>	X ¹	X	X
import existing PDF documents	all PDI functions	-	X	X
Query information from PDF with pCOS	all pCOS functions	-	X	X
Fill Blocks with variable data	<code>PDF_fill_*block()</code>	-	-	X
Create Blocks programmatically	<code>PDF_poca_new()</code> : option usage=blocks <code>PDF_begin/end_page_ext()</code> : option blocks	-	-	X
Copy Blocks to generated output	<code>PDF_process_pdi()</code> : option action=copyblock or action=copyallblocks	-	-	X
interactively create PDFlib Blocks for use with PPS	PDFlib Block Plugin for Acrobat	-	-	X

1. Not available in PDFlib source code packages since PDI is required internally for this feature



2 PDFlib Language Bindings

Note It is strongly recommended to take a look at the starter examples which are included in all PDFlib packages. They provide a convenient starting point for your application development and cover many important aspects of PDFlib programming.

2.1 C Binding

PDFlib is written in C with some C++ modules. In order to use the PDFlib C binding, you can use a static or shared library (DLL on Windows and MVS), and you need the central PDFlib include file *pdflib.h* for inclusion in your PDFlib client source modules. Alternatively, *pdflibdl.h* can be used for dynamically loading the PDFlib DLL at runtime (see next section for details).

Note Applications which use the PDFlib binding for C must be linked with a C++ linker since PDFlib includes parts which are implemented in C++. Using a C linker may result in unresolved externals unless the application is explicitly linked against the C++ runtime libraries.

Error handling in C. PDFlib supports structured exception handling with try/catch clauses. This allows C and C++ clients to catch exceptions which are thrown by PDFlib, and react on the exception in an adequate way. In the catch clause the client will have access to a string describing the exact nature of the problem, a unique exception number, and the name of the PDFlib API method which threw the exception. The general structure of a PDFlib C client program with exception handling looks as follows:

```
PDF *p;

/* create a new PDFlib object */
if ((p = PDF_new()) == (PDF *) 0)
{
    printf("Couldn't create PDFlib object (out of memory)\n");
    return(2);
}
PDF_TRY(p)
{
    ...PDFlib instructions...
}
PDF_CATCH(p)
{
    printf("PDFlib exception occurred:\n");
    printf("[%d] %s: %s\n",
        PDF_get_errnum(p), PDF_get_apiname(p), PDF_get_errmsg(p));
    PDF_delete(p);
    return(2);
}

PDF_delete(p);
```

PDF_TRY/PDF_CATCH are implemented as tricky preprocessor macros. Accidentally omitting one of these will result in compiler error messages which may be difficult to comprehend. Make sure to use the macros exactly as shown above, with no additional code between the *TRY* and *CATCH* clauses (except *PDF_CATCH()*).

An important task of the catch clause is to clean up PDFlib internals using `PDF_delete()` and the pointer to the PDFlib object. `PDF_delete()` will also close the output file if necessary. After an exception the PDF document cannot be used, and will be left in an incomplete and inconsistent state. Obviously, the appropriate action when an exception occurs is application-specific.

For C and C++ clients which do not catch exceptions, the default action upon exceptions is to issue an appropriate message on the standard error channel and exit. The PDF output file is left in an incomplete state! Since this is not adequate for a library routine, for serious PDFlib projects it is strongly advised to leverage PDFlib's exception handling facilities. A user-defined catch clause may, for example, present the error message in a GUI dialog box, and take other measures instead of aborting.

Volatile variables. Special care must be taken regarding variables that are used in both the `PDF_TRY()` and the `PDF_CATCH()` blocks. Since the compiler doesn't know about the control transfer from one block to the other, it might produce inappropriate code (e.g., register variable optimizations) in this situation. Fortunately, there is a simple rule to avoid these problems:

Note Variables used in both the `PDF_TRY()` and `PDF_CATCH()` blocks should be declared volatile.

Using the `volatile` keyword signals to the compiler that it must not apply (potentially dangerous) optimizations to the variable.

Nesting try/catch blocks and rethrowing exceptions. `PDF_TRY()` blocks may be nested to an arbitrary depth. In the case of nested error handling, the inner catch block can activate the outer catch block by re-throwing the exception:

```
PDF_TRY(p)                /* outer try block */
{
    /* ... */

    PDF_TRY(p)             /* inner try block */
    {
        /* ... */
    }
    PDF_CATCH(p)           /* inner catch block */
    {
        /* error cleanup */
        PDF_RETHROW(p);
    }
    /* ... */
}
PDF_CATCH(p)              /* outer catch block */
{
    /* more error cleanup */
    PDF_delete(p);
}
```

The `PDF_RETHROW()` invocation in the inner error handler transfers program execution to the first statement of the outer `PDF_CATCH()` block immediately.

Prematurely exiting a try block. If a `PDF_TRY()` block is left – e.g., by means of a return statement –, thus bypassing the invocation of the corresponding `PDF_CATCH()` macro,

the `PDF_EXIT_TRY()` macro must be used to inform the exception machinery. No other library function must be called between this macro and the end of the try block:

```
PDF_TRY(p)
{
    /* ... */

    if (error_condition)
    {
        PDF_EXIT_TRY(p);
        return -1;
    }
}
PDF_CATCH(p)
{
    /* error cleanup */
    PDF_RETHROW(p);
}
```

Using PDFlib as a DLL loaded at runtime. While most clients will use PDFlib as a statically bound library or a dynamic library which is bound at link time, you can also load the PDFlib DLL at runtime and dynamically fetch pointers to all API methods. This is especially useful to load the PDFlib DLL only on demand. PDFlib supports a special mechanism to facilitate this dynamic usage. It works according to the following rules (this is demonstrated in the *hellodl.c* sample):

- ▶ Include *pdflibdl.h* instead of *pdflib.h*.
- ▶ Use *PDF_new_dl()* and *PDF_delete_dl()* instead of *PDF_new()* and *PDF_delete()*.
- ▶ Use *PDF_TRY_DL()* and *PDF_CATCH_DL()* instead of *PDF_TRY()* and *PDF_CATCH()*.
- ▶ Use function pointers for all other PDFlib calls.
- ▶ *PDF_get_opaque()* must not be used.
- ▶ Compile the auxiliary module *pdflibdl.c* and link your application against it.

In order to create a shared library from the static library on Linux use the following commands:

```
mkdir tmp
cd tmp
ar x ../libpdf.a
g++ -shared -o libpdf.so *
```

This results in a shared library which no longer requires the application to be linked against the C++ runtime library.

Unicode strings in C. You can initialize strings in the following ways:

- ▶ Supply the corresponding surrogate pair (two UTF-16 code units) as four bytes in big-endian byte order:

```
PDF_fit_textline(p, "\x3C\xD8\xCF\xDC", 4, 100, 100,
"fontname=EmojiOneColor textformat=utf16be fontsize=50");
```

Alternatively you can supply the bytes in little-endian byte order:

```
PDF_fit_textline(p, "\xD8\x3C\xDC\xCF", 4, 100, 100,
"fontname=EmojiOneColor textformat=utf16le fontsize=50");
```

- ▶ PDFlib character references (see Section 5.4.2, »Character References«, page 106) are independent of the development environment and can be enabled with the *charref* option, e.g.

```
PDF_fit_textline(p, "&#x1F0CF;", 0,
100, 100, "charref=true fontname=EmojiOneColor fontsize=50");
```

2.2 C++ Binding

An object-oriented wrapper for C++ is available for PDFlib clients. It requires the *pdflib.hpp* header file which in turn includes *pdflib.h*. The template-driven PDFlib interface supports different string types for passing data to and from PDFlib.

Default interface based on *wstring*. Strings of the C++ type *std::wstring* are used for all strings in the API. This is the default string type unless one of the types discussed below is selected.

The default interface assumes that all strings passed to and from PDFlib methods are native *wstrings*. Depending on the size of the *wchar_t* data type, *wstrings* are assumed to contain Unicode strings encoded in UTF-16 (2-byte characters) or UTF-32 (4-byte characters). Literal strings in the source code must be prefixed with *L* to designate wide strings. You can initialize *wstring* literals in the following ways:

- ▶ Supply the Unicode code point with the prefix `\u` and four hexadecimal digits or the prefix `\U` and eight hexadecimal digits where the first three digits are zero for all valid characters (surrogate pairs are not supported):

```
wstring unistring = L"\U0001F0CF";
```

- ▶ Supply the Unicode character directly in the source code and store the source file in a suitable Unicode encoding such as UTF-8 or UTF-16. The compiler must be informed about the source encoding. For example, UTF-8 is the default encoding in GCC for source files, while other encodings can be specified with a command-line option, e.g. *-finput-charset=UTF-16LE*. An option example for Visual Studio is */source-charset:utf-8*.
- ▶ PDFlib character references (see Section 5.4.2, »Character References«, page 106) are independent of the development environment and can be enabled with the *charref* option, e.g.

```
p.fit_textline(L"&#x1F0CF;",
100, 100, "charref=true fontname=EmojiOneColor fontsize=50");
```

Use of the default *std::wstring* interface is demonstrated in the *hello* sample and all other C++ samples.

Note The *wstring* interface cannot be used on IBM Z systems. Use the *u16string/u32string* interfaces instead (see below).

Optional UTF-8/UTF-16/UTF-32 interfaces based on *u8string/u16string/u32string*. The *std::u16string* and *std::u32string* types require at least C++11; the *std::u8string* type requires C++20. These interfaces store Unicode characters in UTF-8, UTF-16 or UTF-32 format. They can be activated by setting one of the *PDFLIB_CPP_U8STRING*, *PDFLIB_CPP_U16STRING* or *PDFLIB_CPP_U32STRING* macros before including *pdflib.hpp*. Their use is demonstrated in the *hello_u8string*, *hello_u16string* and *hello_u32string* samples.

Optional UTF-8 interface based on string. All C++ versions support the `std::string` data type. This interface can be activated by setting the `PDFLIB_CPP_STRING` macro before including `pdflib.hpp`. Unlike the `u8string` interface there is no UTF-8 compiler support with this string type. The client code is responsible for properly filling and interpreting strings with UTF-8 sequences.

Use of the UTF-8-based `std::string` interface is demonstrated in the `hello_string` sample.

Error handling in C++. PDFlib API methods throw a C++ exception in case of an error. These exceptions must be caught in the client code by using C++ `try/catch` clauses. In order to provide extended error information the PDFlib class provides a public `PDFlib::Exception` class which exposes methods for retrieving the detailed error message, the exception number, and the name of the API method which threw the exception.

Native C++ exceptions thrown by PDFlib routines behave as expected. The following code fragment catches exceptions thrown by PDFlib:

```
try {
    PDFlib p;
    ...PDFlib instructions...
} catch (PDFlib::Exception &ex) {
    wcerr << L"PDFlib exception occurred: " << endl
    << L "[" << ex.get_errnum() << L"] " << ex.get_apiname()
    << L": " << ex.get_errmsg() << endl;
}
```

Using PDFlib as a DLL loaded at runtime. The C++ binding allows you to dynamically attach PDFlib to your application at runtime (see »Using PDFlib as a DLL loaded at runtime«, page 29). Dynamic loading can be enabled as follows when compiling the application module which includes `pdflib.hpp`:

```
#define PDFCPP_DL 1
```

In addition you must compile the auxiliary module `pdflibdl.c` and link your application against the resulting object file. Since the details of dynamic loading are hidden in the PDFlib object it does not affect the C++ API: all method calls look the same regardless of whether or not dynamic loading is enabled.

2.3 Java Binding

Installing the PDFlib Java Edition. For the PDFlib binding to work, the Java VM must have access to the PDFlib Java wrapper and the PDFlib Java package. PDFlib is organized as a Java package with the following package name:

```
com.pdflib.pdflib
```

This package is available in the `pdflib.jar` file and contains the classes `pdflib` and `PDFlib-Exception`. In order to supply this package to your application, you must add `pdflib.jar` to your `CLASSPATH` environment variable, add the option `-classpath pdflib.jar` in your calls to the Java compiler and runtime, or perform equivalent steps in your Java IDE. In the JDK you can configure the Java VM to search for native libraries in a given directory by setting the `java.library.path` property to the name of the directory, e.g.

```
java -Djava.library.path=. hello
```

You can check the value of this property as follows:

```
System.out.println(System.getProperty("java.library.path"));
```

In addition, the following platform-dependent steps must be performed:

- ▶ Unix: the library *libpdflic_java.so* (on macOS: *libpdflic_java.jnilib*) must be placed in one of the default locations for shared libraries, or in an appropriately configured directory.
- ▶ Windows: the library *pdflic_java.dll* must be placed in the Windows system directory, or a directory which is listed in the PATH environment variable.

Error handling in Java. The Java binding translates PDFlic errors to native Java exceptions. In case of an exception PDFlic will throw a native Java exception of the class *PDFlicException*. The Java exceptions can be dealt with by the usual try/catch technique:

```
try {
    p = new pdflic();
    ...PDFlic instructions...
} catch (PDFlicException e) {
    System.err.println("PDFlic exception occurred:");
    System.err.println "[" + e.get_errnum() + " ] " + e.get_apiname() +
        ": " + e.get_errmsg());
} catch (Exception e) {
    System.err.println(e);
} finally {
    if (p != null) {
        p.delete();           /* delete the PDFlic object */
    }
}
```

Since PDFlic declares appropriate *throws* clauses, client code must either catch all possible PDFlic exceptions, or declare those itself.

Unicode strings in Java. A *String* object in Java represents strings in UTF-16 format. This means that Unicode values outside the BMP, i.e. > U+FFFF, must be represented as surrogate pairs. Unicode code points outside the BMP cannot be represented in a single *char* or *Character*. You can initialize Strings in the following ways:

- ▶ Supply the Unicode code point as *int* value and convert it to a *String*:

```
String unistring = new String(Character.toChars(0x1F0CF));
```

- ▶ Supply the corresponding surrogate pair (two UTF-16 code units):

```
String unistring = new String("\uD83C\uDCCF");
```

- ▶ Supply the Unicode character directly in the source code and store the source file in a suitable Unicode encoding such as UTF-8 or UTF-16. The compiler must be informed about the source encoding, either directly with the *javac* option *encoding="utf8"* or by appropriate settings in the IDE.

- ▶ PDFlib character references (see Section 5.4.2, »Character References«, page 106) are independent of the development environment and can be enabled with the *charref* option, e.g.

```
p.fit_textline("&#x1F0CF;",
100, 100, "charref=true fontname=EmojiOneColor fontsize=50");
```

Javadoc documentation for PDFlib. The PDFlib package contains Javadoc documentation for PDFlib. The Javadoc contains only abbreviated descriptions of all PDFlib API methods; please refer to the PDFlib API Reference for more details.

In order to configure Javadoc for PDFlib in Eclipse proceed as follows:

- ▶ In the Package Explorer right-click on the Java project and select *Javadoc Location*.
- ▶ Click on *Browse...* and select the path where the Javadoc (which is part of the PDFlib package) is located.

After these steps you can browse the Javadoc for PDFlib, e.g. with the *Java Browsing* perspective or via the *Help* menu.

2.4 .NET Binding

Note The .NET binding is delivered as a universal package for all supported platforms. However, it still requires a platform-specific license key which cannot be transferred across platforms.

Table 2.1 describes general aspects of the .NET binding.

Table 2.1 General properties of the .NET binding

Topic	Implementation of the .NET binding
.NET support	.NET Standard 2.0 and above, e.g. .NET Core 2/3, .NET 5/6
.NET Framework support	.NET Framework 4.6.1 and above
target platforms	Windows, Linux, macOS; see system-requirements.txt for exact list of platform/architecture combinations
download package	universal zip package for all supported platforms
package contents	NuGet package with managed and native libraries, documentation, and samples
implementation	C# assembly PDFlib_dotnet.dll with managed code and platform-specific shared library PDFlib_dotnet_native.[dll so dylib] with native code
.NET integration	C# Interop via explicit PInvoke
versioning scheme	The versioning scheme conforms to .NET rules. The .NET version numbers are visible e.g. in the NuGet cache and .csproj project files. They differ from the product's major and minor release numbers. A mapping between both versioning schemes can be found in compatibility.txt.

Installing PDFlib.NET. The product is supplied as NuGet package where the package name includes the version number. Internal the packages for .NET use semantic version according to .NET rules. As a result, the internal version numbers differ from the regular product version number. A mapping table for both numbering schemes is available in *compatibility.txt*.

The NuGet package can be installed locally using any of the following methods:

- ▶ The *dotnet* command-line tool (all platforms). This method is detailed in the next section.
- ▶ Visual Studio's Package Manager UI (Windows and macOS)
- ▶ Visual Studio's Package Manager Console (Windows)
- ▶ The *nuget* command-line tool (all platforms)

The project files for the supplied samples are prepared for target framework .NET 6.0 (target framework moniker TFM=*net6.0*).

Installing PDFlib for .NET with the dotnet command-line tool. We describe the installation, configuration and build process with the *dotnet* utility, using the supplied *hello* project as an example:

- ▶ Unpack the product package in a directory of your choice.
- ▶ In a command shell *cd* to the *hello* project directory:

```
cd <installdir>\bind\dotnet\csharp\hello
```

- ▶ (This step is not required for the supplied samples which reference the package with a local *NuGet.Config* file) Copy the NuGet package to the application's project directory:

```
<installdir>/bind/dotnet/PDFlib_dotnet.X.Y.Z.nupkg
```

- ▶ (This step is not required for the supplied samples which already contain a reference to PDFlib) Enter the following command with the appropriate version number (the version number can be found in the name of the *.nupkg* file):

```
dotnet add package PDFlib_dotnet.X.Y.Z
```

This command adds a PDFlib reference to the *.csproj* project file. It also installs PDFlib in the local NuGet package cache if it is not yet present, e.g.

```
~/nuget/packages/pdflib_dotnet/X.Y.Z
```

Because of this caching you must copy the **.nupkg* only for the first project. Subsequent projects don't require the package file since it is taken from the cache.

- ▶ Now you can build and run the *hello* project to test it:

```
dotnet build
dotnet run
```

As a result you will find the generated *hello.pdf* output document in the application directory.

Full examples with ready-to-use configuration are included in all packages. Once the .NET binding is properly referenced you can use the *PDFlib_dotnet.PDFlib* and *PDFlib_dotnet.PDFlibException* classes.

Installing PDFlib.NET in a .NET Framework project. Proceed as follows to add PDFlib.NET to a .NET Framework 4.x project with Visual Studio:

- ▶ Add a reference to *PDFlib_dotnet.x.y.z.nupkg* with the NuGet Package manager. As a result a reference to *PDFlib_dotnet* is added to your project as well as a *packages.config* management file.
- ▶ In the context menu of *packages.config* choose *Migrate packages.config to Package-Reference...* to convert the reference (which initially points only to *PDFlib_dotnet.dll*) to a NuGet package reference.

- ▶ In the *Configuration Manager* create a *New Solution Platform* with the name *x86* (32-bit) or *x64* (64-bit). Do not use the setting *Any CPU* since it won't work.
- ▶ If you plan to use the built-in IIS the selected Bitness must match the platform; set *x86* or *x64* appropriately via *Project, Properties, Web, Servers*.
- ▶ If you change the configuration you must rebuild the project to ensure that the correct version of *PDFlib_dotnet_native.dll* is copied to the *bin* directory.

Error handling in .NET. The .NET binding supports .NET exceptions and will throw an exception with a detailed error message when a runtime problem occurs. The client is responsible for catching such an exception and properly reacting on it. Otherwise the .NET framework will catch the exception and usually terminate the application.

In order to convey exception-related information PDFlib defines its own exception class *PDFlib_dotnet.PDFlibException* with the members *get_errnum*, *get_errmsg*, and *get_apiname*. PDFlib implements the *IDisposable* interface which means that clients can call the *Dispose()* method for cleanup.

Client code can handle .NET exceptions thrown by PDFlib with the usual *try...catch* construct:

```
try {
    p = new PDFlib();
    ...PDFlib instructions...
} catch (PDFlibException e) {
    // caught exception thrown by PDFlib
    Console.WriteLine("PDFlib exception occurred:\n");
    Console.WriteLine("[{0}] {1}: {2}\n",
        e.get_errnum(), e.get_apiname(), e.get_errmsg());
} finally {
    if (p != null) {
        p.Dispose();
    }
}
```

Unicode strings in C#. A *String* object in C# represents strings in UTF-16 format. This means that Unicode values outside the BMP, i.e. > U+FFFF, must be represented as surrogate pairs. Unicode code points outside the BMP cannot be represented in a single *char* or *Character*. You can initialize Strings in the following ways:

- ▶ Supply the Unicode code point as *int* value and convert it to a *String*:

```
String unistring = Char.ConvertFromUtf32(0x1F0CF);
```

- ▶ Supply the corresponding surrogate pair (two UTF-16 code units):

```
String unistring = "\uD83C\uDCCF";
```

- ▶ Supply the Unicode character directly in the source code and store the source file in a suitable Unicode encoding such as UTF-8 or UTF-16. The compiler must be informed about the source encoding, either directly or by appropriate settings in the IDE.
- ▶ PDFlib character references (see Section 5.4.2, »Character References«, page 106) are independent of the development environment and can be enabled with the *charref* option, e.g.

```
p.fit_textline("&#x1F0CF;",
    100, 100, "charref=true fontname=EmojiOneColor fontsize=50");
```

2.5 Objective-C Binding

Although the C and C++ language bindings can be used with Objective-C, a genuine language binding for Objective-C is also available. The *PDFlib* framework is available for use on macOS and iOS.

The frameworks contains language bindings for C, C++, and Objective-C.

Installing the PDFlib Edition for Objective-C. In order to use PDFlib in your application you must copy *PDFlib.framework* to the directory */Library/Frameworks*. Installing the framework in a different location is possible, but requires use of Apple's *install_name_tool* which is not described here. The *PDFlib_objc.h* header file with PDFlib method declarations must be imported in the application source code:

```
#import "PDFlib/PDFlib_objc.h"
```

In order to embed the PDFlib framework in an app XCode's code signing expects a framework with the version number *A* while PDFlib products use numeric version numbers. In order to get around this you can create an appropriately named framework folder as follows:

```
cd PDFlib.framework/Versions
mv 10.x.y A
rm Current
ln -s A Current
```

Parameter naming conventions. For PDFlib method calls you must supply parameters according to the following conventions:

- ▶ The value of the first parameter is provided directly after the method name, separated by a colon character.
- ▶ For each subsequent parameter the parameter's name and its value (again separated from each other by a colon character) must be provided. The parameter names can be found in the PDFlib API Reference or in *PDFlib_objc.h*.

For example, the following line in the PDFlib API Reference:

```
void begin_page_ext(double width, double height, String optlist)
```

corresponds to the following Objective-C method:

```
- (void) begin_page_ext: (double) width height: (double) height optlist: (NSString *) optlist;
```

This means your application must make a call similar to the following:

```
[pdflib begin_page_ext:595.0 height:842.0 optlist:@""];
```

XCode Code Sense for code completion can be used with the PDFlib framework.

Error handling in Objective-C. The Objective-C binding translates PDFlib errors to native Objective-C exceptions. In case of a runtime problem PDFlib throws a native Objective-C exception of the class *PDFlibException*. These exceptions can be handled with the usual *try/catch* mechanism:

```
PDFlib p = NULL;
```

```

@try {
    PDFlib *pdflib = [[PDFlib alloc] init];
    ...PDFlib instructions...
}
@catch (PDFlibException *ex) {

    NSLog(@"Error %ld in %@(): %@\n",
          [ex get_errnum], [ex get_apiname], [ex get_errmsg]);
}
@catch (NSException *ex) {

    NSLog(@"%@: %@", [ex name], [ex reason]);
}
@finally {
    if (pdflib)
        [pdflib release];
}

```

In addition to the *get_errmsg* method you can also use the *reason* field of the exception object to retrieve the error message.

2.6 Perl Binding

The PDFlib wrapper for Perl consists of a C wrapper file and two Perl package modules, one for providing a Perl equivalent for each PDFlib API method and another one for the PDFlib object. The C module is used to build a shared library which the Perl interpreter loads at runtime, with some help from the package file. Perl scripts refer to the shared library module via a *use* statement.

Installing the PDFlib Perl Edition. The Perl extension mechanism loads shared libraries at runtime through the DynaLoader module. The Perl executable must have been compiled with support for shared libraries (this is true for the majority of Perl configurations).

For the PDFlib binding to work, the Perl interpreter must access the PDFlib Perl wrapper and the modules *pdflib_pl.pm* and *PDFlib/PDFlib.pm*. In addition to the platform-specific methods described below you can add a directory to Perl's *@INC* module search path using the *-I* command line option:

```
perl -I/path/to/pdflib hello.pl
```

Unix. Perl will search *pdflib_pl.so* (on macOS: *pdflib_pl.bundle*), *pdflib_pl.pm* and *PDFlib/PDFlib.pm* in the current directory or the directory printed by the following Perl command:

```
perl -e 'use Config; print $Config{sitearchexp};'
```

Perl will also search the subdirectory *auto/pdflib_pl*. Typical output of the above command looks like

```
/usr/lib/perl5/site_perl/5.32/i686-linux
```

Windows. The DLL *pdflib_pl.dll* and the modules *pdflib_pl.pm* and *PDFlib/PDFlib.pm* will be searched in the current directory or the directory printed by the following Perl command:

```
perl -e "use Config; print $Config{sitearchexp};"
```

Typical output of the above command looks like

```
C:\Program Files\Perl5.32\site\lib
```

Error Handling in Perl. The Perl binding translates PDFlib errors to native Perl exceptions. The Perl exceptions can be dealt with by applying the appropriate language constructs, i.e., by bracketing critical sections:

```
eval {
    my $p = new PDFlib::PDFlib;
    ...PDFlib instructions...
};
if ($?) {
    die("$0: PDFlib Exception occurred:\n$?");
}
```

Unicode strings in Perl. You can initialize strings in the following ways;

- ▶ Supply the Unicode code point as hexadecimal value (surrogate pairs are not supported):

```
$unistring="\x{1F0CF}";
```

- ▶ Supply a Unicode character name:

```
$unistring = "\N{PLAYING CARD BLACK JOKER}";
```

- ▶ Supply the Unicode character directly in the source code and store the source file in a suitable Unicode encoding. UTF-8 source files must include the *use utf8* pragma. UTF-16 source files must start with a BOM.

- ▶ PDFlib character references (see Section 5.4.2, »Character References«, page 106) are independent of the development environment and can be enabled with the *charref* option, e.g.

```
p.fit_textline("&#x1F0CF;",
100, 100, "charref=true fontname=EmojiOneColor fontsize=50");
```

2.7 PHP Binding

Installing the PDFlib PHP Edition. Detailed information about the various flavors and options for using PDFlib with PHP can be found in the *PDFlib-in-PHP-HowTo.pdf* document which is contained in the distribution packages and also available on the PDFlib Web site.

You must configure PHP so that it knows about the PDFlib library. Add the following line in *php.ini*:

```
extension=php_pdflib
```

PHP searches the library in the directory specified in the *extension_dir* variable in *php.ini* on Unix and additionally in the standard system directories on Windows. You can test which version of the PHP PDFlib binding you have installed with the following one-line PHP script:

```
<?phpinfo()?>
```

This will display a long info page about your current PHP configuration. On this page check the section titled *PDFlib*.

Modified error return for PDFlib functions in PHP. Since PHP uses the convention of returning the value 0 (FALSE) when an error occurs within a method, all PDFlib API methods have been adjusted to return 0 instead of -1 in case of an error. This difference is noted in the method descriptions in the PDFlib API Reference. However, take care when reading Java code fragments, since they use the usual PDFlib convention of returning -1 in case of an error.

File name handling in PHP. Unqualified file names (without any path component) and relative file names for PDF, image, font and other disk files are handled differently in Unix and Windows versions of PHP:

- ▶ PHP on Unix systems will find files without any path component in the directory where the script is located.
- ▶ PHP on Windows will find files without any path component only in the directory where the PHP DLL is located. PDFlib expects UTF-8-encoded file names, while PHP's *dirname()* method on Windows uses the current system codepage. In this case you must convert directory or file names to UTF-8; see Section 5.2.2, »Language Bindings with UTF-8 Support«, page 97, for details. Example:

```
$searchpath = dirname(dirname(__FILE__)).'/data';  
$searchpath = $p->convert_to_unicode("auto", $searchpath, "outputformat=utf8");
```

In order to provide platform-independent file name handling the use of PDFlib's *SearchPath* facility is strongly recommended (see Section 3.1.4, »Resource Configuration and File Search«, page 51).

Exception handling in PHP. Since PHP supports structured exception handling, PDFlib exceptions will be propagated as PHP exceptions. PDFlib will throw an exception of the class *PDFlibException*, which is derived from PHP's standard *Exception* class. You can use the standard *try/catch* technique to deal with PDFlib exceptions:

```
try {  
    $p = new PDFlib();  
    ...PDFlib instructions...  
} catch (PDFlibException $e) {  
    print "PDFlib exception occurred:\n";  
    print "[" . $e->get_errnum() . "] " . $e->get_apiname() . ": "  
        $e->get_errmsg() . "\n";  
}  
catch (Throwable $e) {  
    die("PHP exception occurred: " . $e->getMessage() . "\n");  
}
```

Developing with Eclipse and Zend Studio. The PHP Development Tools (PDT) support PHP development with Eclipse and Zend Studio. PDT can be configured to support context-sensitive help with the steps outlined below.

Add PDFlib to the Eclipse preferences so that it will be known to all PHP projects:

- ▶ Select *Window, Preferences, PHP, Source Paths, Libraries, New...* to launch a wizard.
- ▶ In *User library name* enter *PDFlib*, click *Add External folder...* and select the folder *bind\php\Eclipse PDT*.

In an existing or new PHP project you can add a reference to the PDFlib library as follows:

- ▶ In the PHP Explorer right-click on the PHP project and select *Include Path, Configure Include Path...*
- ▶ Go to the *Libraries* tab, click *Add Library...*, and select *User Library, PDFlib*.

After these steps you can explore the list of PDFlib methods under the *PHP Include Path/PDFlib/PDFlib* node in the PHP Explorer view. When writing new PHP code Eclipse will assist with code completion and context-sensitive help for all PDFlib methods.

Unicode strings in PHP. You can initialize strings in the following ways:

- ▶ Supply the Unicode code point in hexadecimal format (surrogate pairs are not supported):

```
$unistring = "\u{1F0CF}";
```

- ▶ Supply the Unicode character directly in the PHP source and store the source file in UTF-8 format.
- ▶ PDFlib character references (see Section 5.4.2, »Character References«, page 106) are independent of the development environment and can be enabled with the *charref* option, e.g.

```
p.fit_textline("&#x1F0CF;",  
100, 100, "charref=true fontname=EmojiOneColor fontsize=50");
```

2.8 Python Binding

Installing the PDFlib Python Edition. The Python extension mechanism works by loading shared libraries at runtime. For the PDFlib binding to work, the Python interpreter must have access to the PDFlib library for Python which will be searched in the directories listed in the PYTHONPATH environment variable. The name of the Python wrapper depends on the platform:

- ▶ Unix and macOS: *pdflib_py.so*
- ▶ Windows: *pdflib_py.pyd*

In addition to the PDFlib library the following files must be available in the same directory where the library sits:

- ▶ *PDFlib/PDFlib.py*
- ▶ *PDFlib/ __init__.py* (only for Python 2.7)

Error Handling in Python. The Python binding throws a *PDFlibException* in case of an error.. The Python exceptions can be dealt with by the usual *try/catch* technique:

```

try:
    p = PDFlib()

    ...PDFlib instructions...

except PDFlibException:
    print("PDFlib exception occurred:")
    print("[%d] %s: %s" % (ex.errnum, ex.apiname, ex.errmsg))

except Exception:
    print(ex)

finally:
    if p:
        p.delete()

```

Unicode strings in Python. You can initialize strings in the following ways:

- ▶ Supply the Unicode code point with eight hexadecimal digits:

```
$unistring = "\U0001F0CF";
```

- ▶ Surrogates can be supplied as follows:

```
unistring = "\uD83C\uDCCF".encode('utf-16', 'surrogatepass').decode('utf-16')
```

- ▶ Supply the Unicode character directly in the Python source and store the source file in UTF-8 format.
- ▶ PDFlib character references (see Section 5.4.2, »Character References«, page 106) are independent of the development environment and can be enabled with the *charref* option, e.g.

```
p.fit_textline("&#x1F0CF;",
100, 100, "charref=true fontname=EmojiOneColor fontsize=50");
```

2.9 RPG Binding

PDFlib provides a */copy* module that defines all prototypes and some useful constants needed to compile ILE-RPG programs with embedded PDFlib functions.

Unicode string handling. Since all functions provided by PDFlib use Unicode strings with variable length as parameters, you have to use the *%UCS2* builtin function to convert a single-byte string to a Unicode string. All strings returned by PDFlib functions are Unicode strings with variable length. Use the *%CHAR* builtin function to convert these Unicode strings to single-byte strings.

Note The *%CHAR* and *%UCS2* functions use the current job's *CCSID* to convert strings from and to Unicode. The examples provided with PDFlib are based on *CCSID 37 (US EBCDIC)*. This codepage can be set with *CHGJOB CCSID(37)*. Some characters in option lists (e.g. { [] }) may not be translated correctly if you run the examples under other codepages.

Since all strings are passed as variable length strings you must not pass the *length* parameters in various functions which expect explicit string lengths (the length of a variable length string is stored in the first two bytes of the string).

Compiling and binding RPG Programs for PDFlib. Using PDFlib functions from RPG requires the compiled PDFLIB and PDFLIB_RPG service programs. To include the PDFlib definitions at compile time you have to specify the name of the */copy* member in the *D* specs of your ILE-RPG program:

```
d/copy QRPGLESRC,PDFLIB
```

If the PDFlib source file library is not on top of your library list you have to specify the library as well:

```
d/copy PDFsrcLib/QRPGLESRC,PDFLIB
```

Before you start compiling your ILE-RPG program you have to create a binding directory that includes the PDFLIB and PDFLIB_RPG service programs shipped with PDFlib. The following example assumes that you want to create a binding directory called PDFLIB in the library PDFLIB:

```
CRTBNDDIR BNDDIR(PDFLIB/PDFLIB) TEXT('PDFlib Binding Directory')
```

After creating the binding directory you need to add the PDFLIB and PDFLIB_RPG service programs to your binding directory. The following example assumes that you want to add the service program PDFLIB in the library PDFLIB to the binding directory created earlier.

```
ADDBNDDIRE BNDDIR(PDFLIB/PDFLIB) OBJ((PDFLIB/PDFLIB *SRVPGM))
```

Now you can compile your program using the *CRTBNDRPG* command (or option 14 in PDM):

```
ADDLIBLE LIB(PDFLIB)  
CRTBNDRPG PGM(PDFLIB/HELLO) SRCFILE(PDFLIB/QRPGLESRC) SRCMBR(*PGM) DFACTGRP(*NO)  
BNDDIR(PDFLIB/PDFLIB)
```

Before executing the generated program you should apply the following command:

```
chgcurdir '/pdflib/pdflib/x.y/x.y.z/bind/rpg'
```

The generated PDF documents can be found in the same directory.

Parameters must be passed to the PDFlib API according to the data types listed in Table 2.2.

Table 2.2 Data types in the RPG binding

API data type	data types in the RPG binding
string data type	Unicode string (use %ucs2)
binary data type	data

Error Handling in RPG. PDFlib clients written in ILE-RPG can use the *monitor/on-error/endmon* error handling mechanism that ILE-RPG provides. Another way to monitor for exceptions is to use the **PSSR* global error handling subroutine in ILE-RPG. If an exception occurs, the job log shows the error number, the function that failed and the reason for the exception. PDFlib sends an escape message to the calling program.

```

c   eval      p=PDF_new
*
c   monitor
*
c   eval      doc=PDF_begin_document(p:%ucs2('/tmp/my.pdf'):docoptlist)
:
:
*   Error Handling
c   on-error
*   Do something with this error
*   don't forget to free the PDFlib object
c   callp     PDF_delete(p)
c   endmon

```

2.10 Ruby Binding

Installing the PDFlib Ruby edition. The Ruby extension mechanism works by loading a shared library at runtime. For the PDFlib binding to work, the Ruby interpreter must have access to the PDFlib extension library for Ruby. This library (on Windows and Unix: *PDFlib.so*; on macOS: *PDFlib.bundle*) will usually be installed in the *site_ruby* branch of the local ruby installation directory, i.e. in a directory with a name similar to the following:

```
/usr/local/lib/ruby/site_ruby/<version>/
```

However, Ruby will search other directories for extensions as well. In order to retrieve a list of these directories you can use the following ruby call:

```
ruby -e "puts $:"
```

This list will usually include the current directory, so for testing purposes you can simply place the PDFlib extension library and the scripts in the same directory.

Error Handling in Ruby. The Ruby binding installs an error handler which translates PDFlib exceptions to native Ruby exceptions. The Ruby exceptions can be dealt with by the usual *rescue* technique:

```

begin
  p = PDFlib.new
  ...PDFlib instructions...
rescue PDFlibException => pe
  print "PDFlib exception occurred:\n"
  print "[" + pe.get_errnum.to_s + "]" + pe.get_apiname + ": " + pe.get_errmsg + "\n"
rescue Exception => e
  print e.backtrace.join("\n") + "\n" + e.to_s
ensure
  p.delete() if p
end

```

Unicode strings in Ruby. You can initialize strings in the following ways:

- Supply the Unicode code point with four, five or six hexadecimal digits (surrogate pairs are not supported):

```
$unistring = "\u{1FOCF}";
```

- ▶ Supply the Unicode character directly in the Ruby source and store the source file in UTF-8 format.
- ▶ PDFlib character references (see Section 5.4.2, »Character References«, page 106) are independent of the development environment and can be enabled with the *charref* option, e.g.

```
p.fit_textline("&#x1F0CF;",
100, 100, "charref=true fontname=EmojiOneColor fontsize=50");
```

Ruby on Rails. Ruby on Rails is an open-source framework which facilitates Web development with Ruby. The PDFlib extension for Ruby can be used with Ruby on Rails. Follow these steps to run the PDFlib examples for Ruby on Rails:

- ▶ Install Ruby and Ruby on Rails.
- ▶ Set up a new controller from the command line:

```
$ rails new pdflibdemo
$ cd pdflibdemo
$ cp <PDFlib dir>/bind/ruby/<version>/PDFlib.so vendor/ # use .so/.dll/.bundle
$ rails generate controller home demo
$ rm public/index.html
```

- ▶ Edit *config/routes.rb*:

```
...
# remember to delete public/index.html
root :to => "home#demo"
```

- ▶ Edit *app/controllers/home_controller.rb* as follows and insert PDFlib code for creating PDF contents. Keep in mind that the PDF output must be generated in memory, i.e. an empty file name must be supplied to *begin_document()*. As a starting point you can use the code in the *hello-rails.rb* sample:

```
class HomeController < ApplicationController
  def demo
    require "PDFlib"
    begin
      p = PDFlib.new
      ...
      ...PDFlib application code, see hello-rails.rb...
      ...
      send_data p.get_buffer(), :filename => "hello.pdf",
      :type => "application/pdf", :disposition => "inline"
      rescue PDFlibException => pe
      # error handling
    end
  end
end
```

- ▶ In order to test your installation start the WEBrick server with the command

```
$ rails server
```

and point your browser to *http://0.0.0.0:3000*. The generated PDF document will be displayed in the browser.

Local PDFlib installation. If you want to use PDFlib only with Ruby on Rails, but cannot install it globally for general use with Ruby, you can install PDFlib locally in the *vendors* directory within the Rails tree. This is particularly useful if you do not have permission to install Ruby extensions for general use, but want to work with PDFlib in Rails nevertheless.



3 Creating PDF Documents

3.1 General PDFlib Programming Aspects

Cookbook Code samples regarding general programming issues can be found in the general category of the PDFlib Cookbook.

3.1.1 Exception Handling

Errors of a certain kind are called exceptions in many languages for good reasons – they are mere exceptions, and are not expected to occur very often during the lifetime of a program. The general strategy is to use conventional error reporting mechanisms (i.e. special error return codes such as -1) for calls which may often fail, and use a special exception mechanism for those rare occasions which don't warrant cluttering the code with conditionals. This is exactly the path that PDFlib goes: Some operations can be expected to go wrong rather frequently, for example:

- ▶ Trying to open an output file for which one doesn't have permission
- ▶ Trying to open an input PDF with a wrong file name
- ▶ Trying to open a corrupt image file

PDFlib signals such errors by returning a special value (usually -1, but 0 in the PHP binding) as documented in the PDFlib API Reference. This error code must be checked by the application developer for all functions which are documented to return -1 on error.

Other events may be considered harmful, but will occur rather infrequently, e.g.

- ▶ running out of virtual memory
- ▶ scope violations (e.g., closing a document before opening it)
- ▶ supplying wrong parameters to PDFlib API methods (e.g., trying to draw a circle with negative radius), or supplying wrong options.

When PDFlib detects such a situation, an exception will be thrown instead of passing a special error return value to the caller. It is important to understand that the generated PDF document cannot be finished when an exception occurred. The only methods which can safely be called after an exception are `PDF_delete()`, `PDF_get_apiname()`, `PDF_get_errnum()`, and `PDF_get_errmsg()`. Calling any other PDFlib method after an exception may lead to unexpected results. The exception contains the following information:

- ▶ A unique error number;
- ▶ The name of the PDFlib API method which caused the exception;
- ▶ A descriptive text containing details of the problem.

Querying the reason of a failed call. As noted above, the generated PDF output document must always be abandoned when an exception occurs. However, when a method reports a non-fatal problem by returning an error value, the document can be continued. The client application is free to continue the document by adjusting the program flow or supplying different data. For example, when a particular font cannot be loaded most clients will give up the document, while others may prefer to work with a different font. In this case it may be desirable to retrieve an error message which describes the problem in more detail. In this situation the functions `PDF_get_errnum()`,

`PDF_get_errmsg()`, and `PDF_get_apiname()` can be called immediately after a failed call, i.e., a call which returned a -1 (in PHP: 0) error value.

Error policies. When PDFlib detects an error condition, it will react according to one of several strategies which can be configured with the `errorpolicy` option. All methods which can return error codes also support the `errorpolicy` option. The following error policies are supported:

- ▶ `errorpolicy=return`: when an error condition is detected the method returns with a -1 (in PHP: 0) error value. The application developer must check the return value to identify problems, and must react on the problem in whatever way is appropriate for the application. This is the recommended approach since it allows a unified approach to error handling.
- ▶ `errorpolicy=exception`: an exception is thrown when an error condition is detected. The output document will be incomplete and unusable and must be discarded (this can be automated with the `removefragments` document option). This setting can be used for lazy programming without any error conditionals at the expense of sacrificing the output document even for nonfatal problems which may be fixable by the application.

The following code fragments demonstrate different strategies with respect to exception handling. The examples try to load a font which may or may not be available.

If `errorpolicy=return` the return value must be checked for an error. If it indicates failure, the reason of the failure can be queried in order to properly deal with the situation:

```
font = p.load_font("MyFontName", "unicode", "errorpolicy=return");
if (font == -1)
{
    /* font handle is invalid; find out what happened. */
    errmsg = p.get_errmsg();
    /* Try a different font or give up */
    ...
}
/* font handle is valid; continue */
```

If `errorpolicy=exception` the document must be abandoned if an error occurs:

```
font = p.load_font("MyFontName", "unicode", "errorpolicy=exception");
/* Unless an exception was thrown the font handle is valid;
 * if an exception occurred, the PDF output cannot be continued
 */
```

Cookbook A full code sample can be found in the *Cookbook* topic `general/error_handling`.

Warnings. Some problem conditions are detected by PDFlib internally, but do not justify interrupting the program flow by throwing an exception. Instead of throwing an exception, a description of the condition is logged (see Section 3.1.2, »Logging«, page 49, for more details about the logging feature). We recommend the following approach with respect to warnings:

- ▶ Enable warning logging in the development phase and carefully study any warning messages in the log file. They may point to potential problems in your code or data, and you should try to understand or eliminate the reason for those warnings.

- ▶ Disable warning logging in the production phase, and re-enable it only in case of problems.

3.1.2 Logging

PDFlib can create a log file which records the following items:

- ▶ all API calls with corresponding parameters and options; note that logging may contain additional API calls issued by the PDFlib language wrapper (e.g. for string conversion).
- ▶ return values and handles returned by API methods;
- ▶ a timestamp for each call;
- ▶ information regarding the use of deprecated API methods and options;
- ▶ warning messages which don't justify an exception, but should be examined during the development phase of an application;
- ▶ details about internal operation which may be useful for investigating support cases.

The contents of a log file may be an important aid for application developers for identifying problems in the program flow. Logging can be enabled at runtime as follows:

```
p.set_option("logging={filename={mylogfile.log}}");
```

Alternatively, logging can be enabled via environment variables. The amount of logged information can be controlled with various logging classes and levels for each class; see PDFlib API Reference for details.

Logging should not normally be enabled in production situations, but only during the development phase and when problems need to be analyzed. PDFlib GmbH support may request log files to discuss user problems.

The following excerpt shows lines from a typical log file with default logging classes:

```
PDF_load_image(p_0x2201c20, "auto", "nesrin.jpg", /*c*/0, "")
[0]

PDF_begin_page_ext(p_0x2201c20, 10.000000, 10.000000, "")
[Begin page #1]
PDF_fit_image(p_0x2201c20, 0, 0.000000, 0.000000, "adjustpage")
PDF_close_image(p_0x2201c20, 0)
PDF_end_page_ext(p_0x2201c20, "")
[End page #1]

PDF_end_document(p_0x2201c20, "")
[Full product name: "PDFlib Personalization Server"]
[Document ID: <C98301CB2D4EAC2972D34CAAAE929021> <C98301CB2D4EAC2972D34CAAAE929021>]
[Size of document: 34842 bytes]
[End document #1]
```

3.1.3 The PDFlib Virtual File System (PVF)

Cookbook A full code sample can be found in the *Cookbook* topic `general/data_from_memory`.

In addition to disk files a facility called *PDFlib Virtual File System (PVF)* allows clients to directly supply data in memory without any disk files involved. This offers performance benefits and can be used for data fetched from a database which does not even exist on

an isolated disk file, as well as other situations where the client already has the required data available in memory as a result of some processing.

PVF is based on the concept of named virtual read-only files which can be used just like regular file names with any API method. They can even be used in UPR configuration files. Virtual file names can be generated in an arbitrary way by the client. Obviously, virtual file names must be chosen such that name clashes with regular disk files are avoided. For this reason a hierarchical naming convention for virtual file names is recommended as follows (*filename* refers to a name chosen by the client which is unique in the respective category). It is also recommended to keep standard file name suffixes:

- ▶ Raster image files: */pvf/image/filename*
- ▶ font outline files (it is recommended to use the actual font name as the base portion of the file name): */pvf/font/filename*
- ▶ ICC profiles: */pvf/iccprofile/filename*
- ▶ PDF documents: */pvf/pdf/filename*

When searching for a named file PDFlib will first check whether the supplied file name refers to a known virtual file, and then try to open the named file on disk.

Lifetime of virtual files. Some functions will immediately consume the data supplied in a virtual file, while others will read only parts of the file, with other fragments being used at a later point in time. For this reason close attention must be paid to the lifetime of virtual files. PDFlib will place an internal lock on every virtual file, and remove the lock only when the contents are no longer needed. Unless the client requested PDFlib to make an immediate copy of the data (using the *copy* option in *PDF_create_pvf()*), the virtual file's contents must only be modified, deleted, or freed by the client when it is no longer locked by PDFlib. PDFlib will automatically delete all virtual files in *PDF_delete()*. However, the actual file contents (the data comprising a virtual file) must always be freed by the client.

Different strategies. PVF supports different approaches with respect to managing the memory required for virtual files. These are governed by the fact that PDFlib may need access to a virtual file's contents after the API call which accepted the virtual file name, but never needs access to the contents after *PDF_end_document()*. Remember that calling *PDF_delete_pvf()* does not free the actual file contents (unless the *copy* option has been supplied), but only the corresponding data structures used for PVF file name administration. This gives rise to the following strategies:

- ▶ Minimize memory usage: it is recommended to call *PDF_delete_pvf()* immediately after the API call which accepted the virtual file name, and another time after *PDF_end_document()*. The second call is required because PDFlib may still need access to the data so that the first call refuses to unlock the virtual file. However, in some cases the first call will already free the data, and the second call doesn't do any harm. The client may free the file contents only when *PDF_delete_pvf()* succeeded.
- ▶ Optimize performance by reusing virtual files: some clients may wish to reuse some data (e.g., font definitions) within various output documents, and avoid multiple create/delete cycles for the same file contents. In this case it is recommended not to call *PDF_delete_pvf()* as long as more PDF output documents using the virtual file will be generated.

- ▶ **Lazy programming:** if memory usage is not a concern the client may elect not to call `PDF_delete_pvf()` at all. In this case PDFlib will internally delete all pending virtual files in `PDF_delete()`.

In all cases the client may free the corresponding data only when `PDF_delete_pvf()` returned successfully, or after `PDF_delete()`.

Creating PDF output in a virtual file. In addition to supplying user data to PDFlib, PVF can also hold the PDF document data generated by PDFlib. This can be achieved by supplying the `createpvf` option to `PDF_begin_document()`. The PVF file name can later be supplied to other PDFlib API methods. This is useful, for example, when generating PDF documents for inclusion in a PDF Portfolio. It is not possible to directly retrieve the PVF data created by PDFlib; use the active or passive in-core PDF generation interface to fetch PDF data from memory (see Section 3.1.5, »Generating PDF Documents in Memory«, page 55).

3.1.4 Resource Configuration and File Search

In most advanced applications PDFlib needs access to resources such as font files, ICC color profiles, etc. In order to make PDFlib’s resource handling platform-independent and customizable, a configuration file can be supplied for describing the available resources along with the names of their corresponding disk files. In addition to a static configuration file, dynamic configuration can be accomplished at runtime by adding resources with `PDF_set_option()`. For the configuration file PDFlib uses a simple text format called *Unix PostScript Resource* (UPR). The UPR file format as used by PDFlib is described below.

With the `enumeratefonts` option PDFlib can be instructed to collect all fonts which are accessible on the search path (see »Automatically creating a font resource list«, page 122).

Resource categories. The resource categories supported by PDFlib are listed in Table 3.1. Most categories map a resource name (to be used in the PDFlib API) to the name of a virtual or disk-based file. Resource categories other than those in Table 3.1 are ignored. Category names are case-insensitive. The values are treated as name strings; they can be encoded in ASCII or UTF-8 (with BOM at the start of a line), or in EBCDIC-UTF-8 on IBM Z. Unicode values may be useful for localized font names with the *HostFont* resource.

Table 3.1 Resource categories supported in PDFlib

<i>category</i>	<i>format</i>	<i>explanation</i>
SearchPath	pathname	<i>relative or absolute path name of directories containing data files</i>
CMap	cmapname=filename	<i>(Deprecated) CMap file for CJK encoding</i>
FontAFM	fontname=filename	<i>(Deprecated) PostScript font metrics file in AFM format</i>
FontPFM	fontname=filename	<i>(Deprecated) PostScript font metrics file in PFM format</i>
FontOutline	fontname=filename	<i>PostScript, TrueType, OpenType, WOFF, or WOFF2 font file</i>
Encoding	encodingname=filename	<i>text file containing an 8-bit encoding or code page table</i>
HostFont	fontname=hostfontname	<i>name of a font installed on the system (usually both font names are identical)</i>

Table 3.1 Resource categories supported in PDFlib

category	format	explanation
FontnameAlias	aliasname=fontname	create an alias for a font which is already known to PDFlib
ICCProfile	profilename=filename	name of an ICC color profile

The UPR file format. UPR files are text files with a very simple structure that can easily be written in a text editor or generated automatically. To start with, let's take a look at some syntactical issues:

- ▶ Lines can have a maximum of 1023 characters.
- ▶ A backslash character '\ ' at the end of a line cancels the line end. This may be used to extend lines.
- ▶ A percent '%' character introduces a comment until the end of the line. Percent characters which are part of the line data (i.e. which do not start a comment) must be protected with a preceding backslash character.
- ▶ Backslash characters in front of a backslash which protects the line end and backslash characters which protect a percent character must be duplicated if they are part of the line data.
- ▶ An isolated period character '.' serves as a section terminator.
- ▶ All entries are case-sensitive.
- ▶ Whitespace is ignored everywhere except in resource names and file names.
- ▶ Resource names and values must not contain any equal character '='.
- ▶ If a resource is defined more than once, the last definition will overwrite earlier definitions.

UPR files consist of the following components:

- ▶ A magic line for identifying the file. It has the following form:

```
PS-Resources-1.0
```

- ▶ An optional section listing all resource categories described in the file. Each line describes one resource category. The list is terminated by a line with a single period character. Available resource categories are described below. If this optional section is not present, a single period character must be present nevertheless.
- ▶ A section for each of the resource categories listed at the beginning of the file. Each section starts with a line showing the resource category, followed by an arbitrary number of lines describing available resources. The list is terminated by a line with a single period character. Each resource data line contains the name of the resource (equal signs have to be quoted). If the resource requires a file name, this name has to be added after an equal sign. The *SearchPath* (see below) will be applied when PDFlib searches for files listed in resource entries.

File search and the SearchPath resource category. PDFlib reads a variety of data items, such as raster images, font outline and metrics information, PDF documents, and ICC color profiles from disk files. In addition to relative or absolute path names you can also use file names without any path specification. The *SearchPath* resource category can be used to specify a list of path names for directories containing the required data files. When PDFlib must open a file it will first use the file name exactly as supplied and try to open the file. If this attempt fails, PDFlib tries to open the file in the directories specified in the *SearchPath* resource category one after another until it succeeds. *SearchPath* en-

tries can be accumulated, and are searched in reverse order (paths set at a later point in time will be searched before earlier ones). This feature can be used to free PDFlib applications from platform-specific file system schemes. You can set search path entries as follows:

```
p.set_option("SearchPath={{/path/to/dir1} {/path/to/dir2}}");
```

The search path can be set multiply, and multiple directory names can be supplied in a single call. It is recommended to use double braces even for a single entry to avoid problems with directory names containing space characters. An empty string list (i.e. `{{}}`) deletes all existing search path entries including the default entries.

In order to disable the search you can use a fully specified path name in the PDFlib functions. Note the following platform-specific features of the *SearchPath* resource category:

On Windows PDFlib initializes the *SearchPath* with entries from the registry. The following registry entries may contain a list of path names separated by a semicolon ';' character. They will be searched in the order provided below:

```
HKLM\SOFTWARE\PDFlib\PDFlib10\10.0.1\SearchPath
HKLM\SOFTWARE\PDFlib\PDFlib10\SearchPath
HKLM\SOFTWARE\PDFlib\SearchPath
```

Default file search paths. On Unix, Linux, macOS and IBM System i systems some directories will be searched for files by default even without specifying any path and directory names. Before searching and reading the UPR file (which may contain additional search paths), the following directories are searched:

```
<rootpath>/PDFlib/PDFlib/10.0/resource/cmap
<rootpath>/PDFlib/PDFlib/10.0/resource/codepage
<rootpath>/PDFlib/PDFlib/10.0/resource/font
<rootpath>/PDFlib/PDFlib/10.0/resource/icc
<rootpath>/PDFlib/PDFlib/10.0
<rootpath>/PDFlib/PDFlib
<rootpath>/PDFlib
```

On Unix, Linux, and macOS *<rootpath>* is first replaced with the HOME directory and then with */usr/local*. On IBM System i *<rootpath>* is empty.

The last of the entries above is especially useful for storing a license file for multiple products.

If the default search paths are not required you can accelerate file operations by disabling them by supplying an empty *searchpath*.

Default file names for license and resource files. By default, the following file names are searched for in the default search path directories:

```
licensekeys.txt          (license file; on MVS: license)
pdflib.upr              (resource file; on MVS: upr)
```

This feature can be used to work with a license file without setting any environment variable or runtime option.

Sample UPR file. The following listing gives an example of a UPR configuration file:

```

PS-Resources-1.0
.
SearchPath
/usr/local/lib/fonts
/users/kurt/my_images
.
FontOutline
ArialMT=Arial.ttf
.
HostFont
Wingdings=Wingdings
.
ICCProfile
highspeedprinter=cmykhigspeed.icc
.

```

Searching for the UPR resource file. If only the built-in resources (e.g., PDF core font, sRGB ICC profile) or system resources (host fonts) are to be used, a UPR configuration file is not required since PDFlib will find all necessary resources without any additional configuration.

If other resources are to be used you can specify such resources via calls to `PDF_set_option()` (see below) or in a UPR resource file. PDFlib reads this file automatically when the first resource is requested. The detailed process is as follows:

- ▶ On Unix systems, macOS and IBM System i some directories are searched by default for license and resource files even without specifying any path and directory names. Before searching and reading the UPR file, the following directories are searched (in this order):

```

<rootpath>/PDFlib/PDFlib/10.0/resource/icc
<rootpath>/PDFlib/PDFlib/10.0/resource/fonts
<rootpath>/PDFlib/PDFlib/10.0/resource/cmap
<rootpath>/PDFlib/PDFlib/10.0
<rootpath>/PDFlib/PDFlib
<rootpath>/PDFlib

```

On Unix systems and macOS `<rootpath>` will first be replaced with `/usr/local` and then with the HOME directory. On IBM System i `<rootpath>` is empty. This feature can be used to work with a license file, UPR file, or resources without setting any environment variables or runtime options.

- ▶ If the environment variable `PDFLIBRESOURCEFILE` is defined PDFlib takes its value as the name of the UPR file to be read. If this file cannot be read an exception will be thrown.
- ▶ If the environment variable `PDFLIBRESOURCEFILE` is not defined PDFlib tries to open a file with the following name:

```

upr (on MVS; a dataset is expected)
pdflib/<version>/fonts/pdflib.upr (on IBM System i)
pdflib.upr (Windows, Unix, and all other systems)

```

If this file cannot be read no exception will be thrown.

- ▶ On Windows PDFlib additionally tries to read the following registry entries in the order provided below:

```

HKLM\Software\PDFlib\PDFlib10\10.0.1\resourcefile
HKLM\Software\PDFlib\PDFlib10\resourcefile

```

```
HKLM\Software\PDFlib\resourcefile
```

The values of these entries will be taken as the name of the resource file to be used. If this file cannot be read an exception will be thrown. Be careful when manually accessing the registry on 64-bit Windows systems: as usual, 64-bit PDFlib binaries will work with the 64-bit view of the Windows registry, while 32-bit PDFlib binaries running on a 64-bit system will work with the 32-bit view of the registry. If you must add registry keys for a 32-bit product manually, make sure to use the 32-bit version of the *regedit* tool. It can be invoked as follows from the *Start* dialog:

```
%systemroot%\syswow64\regedit
```

- ▶ The client can force PDFlib to read a resource file at runtime by explicitly setting the *resourcefile* option:

```
p.set_option("resourcefile={/path/to/pdflib.upr}");
```

This call can be repeated arbitrarily often; the resource entries will be accumulated.

Configuring resources at runtime. In addition to using a UPR file for the configuration, it is also possible to directly configure individual resources within the source code via *PDF_set_option()*. This method takes a category name and a corresponding resource entry as it would appear in the respective section of this category in a UPR resource file, for example:

```
p.set_option("FontOutline={Foobar-Bold=foobb.otf}");
```

Note Font configuration is discussed in more detail in Section 6.3.4, »Searching for Fonts«, page 122.

Querying resource values. In addition to setting resource entries you can query values using *PDF_get_option()*. Specify the category name as key and the number of the resource (starting at 1) as option. For example, the following call:

```
idx = p.get_option("SearchPath", "resourcenumber=" + n);  
sp = p.get_string(idx, "");
```

retrieves the *n*-th entry in the SearchPath list. If *n* is larger than the number of available entries for the requested category an empty string will be returned. The returned string is valid until the next call to any API method.

3.1.5 Generating PDF Documents in Memory

In addition to generating PDF documents on a file, PDFlib can also be instructed to generate the PDF directly in memory (*in-core*). This technique offers performance benefits since no disk-based I/O is involved, and the PDF document can, for example, directly be streamed via HTTP. Webmasters will be especially happy to hear that their server will not be cluttered with temporary PDF files.

You may, at your option, periodically collect partial data (e.g., every time a page has been finished), or fetch the complete PDF document in a single chunk at the end (after *PDF_end_document()*). Interleaving production and consumption of the PDF data has several advantages. Firstly, since not all data must be kept in memory, the memory requirements are reduced. Secondly, such a scheme can boost performance since the first chunk of data can be transmitted over a slow link while the next chunk is still being

generated. However, the total length of the generated data will only be known when the complete document is finished.

You can use the *createpdf* option to create PDF data in memory and subsequently pass it to PDFlib without writing a disk file (see »Creating PDF output in a virtual file«, page 51).

The active in-core PDF generation interface. In order to generate PDF data in memory, simply supply an empty filename to *PDF_begin_document()*, and retrieve the data with *PDF_get_buffer()*:

```
p.begin_document("", "");
...create document...
p.end_document("");

buf = p.get_buffer();
... use the PDF data contained in the buffer ...
p.delete();
```

Note The PDF data in the buffer must be treated as binary data.

This is considered »active« mode since the client decides when he wishes to fetch the buffer contents. Active mode is available for all supported language bindings.

Note C and C++ clients must not free the returned buffer.

The passive in-core PDF generation interface. In »passive« mode, which is only available in the C and C++ language bindings, the user installs (via *PDF_open_document_callback()*) a callback function which will be called at unpredictable times by PDFlib whenever PDF data is waiting to be consumed. Timing and buffer size constraints related to flushing (transferring the PDF data from the library to the client) can be configured by the client in order to provide for maximum flexibility. Depending on the environment, it may be advantageous to fetch the complete PDF document at once, in multiple chunks, or in many small segments in order to prevent PDFlib from increasing the internal document buffer. The flushing strategy can be set using the *flush* option of *PDF_open_document_callback()*.

3.1.6 Maximum Size of PDF Documents and other Limits

Size of PDF documents. Although most users won't see any need for PDF documents in the range of Gigabytes, some enterprise applications must create or process documents containing a large number of, say, invoices or statements. While PDFlib itself does not impose any limits on the size of the generated documents, there are several restrictions mandated by the PDF Reference and some PDF standards:

- ▶ 10 GB file size limit: PDF 1.4 documents are limited by the cross-reference table to 10 decimal digits and therefore $10^{10}-1$ bytes, which equates to roughly 9.3 GB. If you plan to create output documents beyond 10 GB you must use PDF 1.5 or above. This version supports cross-reference streams which are no longer subject to the 10-digits limit and therefore allow creation of PDF documents beyond 10 GB.
- ▶ Number of objects: while the object count in a document is not limited by PDF in general, the PDF/A, PDF/X-4 and PDF/X-5 standards limit the number of indirect objects in a document to 8.388.607. If a document requires objects beyond this limit PDFlib will throw an exception in PDF/A, PDF/X-4 and PDF/X-5 mode. In other modes

documents with more objects can always be created. This check can be disabled with the document option *limitcheck=false*.

The number of objects in PDF depends on the complexity of the page contents, number of interactive elements, etc. Since typical high-volume documents with simple contents require ca. 4-10 objects per page on average, documents with ca. 1-2 million pages can be created without exceeding the object limit mandated by the standards.

PDF limits. PDFlib imposes limits on certain entities in order to create PDF output which conforms to the limitations imposed by the PDF Reference, Acrobat, or some PDF standard. These limits are documented below.

The following limits are enforced by suitably modifying the values:

- ▶ Smallest absolute floating point value in PDF: 0.000015. Numbers with a smaller absolute value are replaced with 0.
- ▶ (PDF 1.4, but not newer PDF versions) Largest absolute value which can be expressed as floating point number in PDF: 32767.0. Numbers with a larger absolute value are replaced with the closest integer.

The PDF format imposes certain restrictions. Exceeding one of the following limits results in an exception:

- ▶ Largest allowed numerical value in PDF: 2.147.483.647
- ▶ Maximum length of hypertext strings: 65535
- ▶ Maximum length of text strings on the page: 32.763 bytes (i.e. 16.381 characters for CID fonts) if *kerneling=false* and *wordspacing=0*; otherwise 4095 characters
- ▶ The following options are limited to a maximum of 8191 list entries: *views*, *namelist*, *polylinelist*, *fieldnamelist*, *itemnamelist*, *itemtextlist*, *children*, *group*
- ▶ Maximum number of indirect objects in PDF/A-1/2/3 and PDF/X-4/5: 8.388.607

3.1.7 Web-Optimized (Linearized) PDF

PDFlib can apply a process called linearization to PDF documents (linearized PDF is also called *Optimized* or *Fast Web View*). Linearization reorganizes the objects within a PDF file and adds supplemental information which can be used for faster access.

While non-linearized PDFs must be fully transferred to the client, a Web server can transfer linearized PDF documents one page at a time using a process called byte-serving. It allows the Web browser to retrieve individual parts of a PDF document separately. The result is that the first page of the document is presented to the user without having to wait for the full document to be downloaded from the server.

The PDF document must be linearized to take advantage of byteserving PDFs. This can be achieved with the *linearize* option in *PDF_begin_document()*:

```
p.begin_document(outfilename, "linearize");
```

3.1.8 Multi-threaded Programming

While PDFlib itself is single-threaded, it can safely be used in multi-threaded applications. In the common situation that a PDFlib object is only used within one thread, no particular multi-threading precautions are necessary. If the same object is used in multiple threads the application must synchronize the threads to make sure that the PDFlib object is not accessed simultaneously by more than one thread. A typical scenario would involve a pool of PDFlib objects where each thread fetches an existing PDFlib object from the pool instead of creating a new one, and returns it to the pool after creating

a document if the object is no longer needed. Using the same PDFlib object in another thread before the output document is finished will rarely provide any advantage and is not recommended.

3.1.9 Using PDFlib on EBCDIC-based Platforms

The operators and structure elements in the PDF file format are based on ASCII which doesn't work well with the EBCDIC-based platforms IBM System i and IBM Z (but not zLinux which is based on ASCII). However, a special mainframe version of PDFlib is available in order to allow mixing of ASCII-based PDF operators and EBCDIC (or other) text output. The EBCDIC-safe version of PDFlib is available for various operating systems and machine architectures.

In order to leverage PDFlib's features on EBCDIC-based platforms the following items are expected to be supplied in EBCDIC text format (more specifically, in code page 037 on IBM System i, and code page 1047 on IBM Z):

- ▶ UPR configuration files
- ▶ encoding and code page files
- ▶ string parameters to PDFlib functions
- ▶ input and output file names
- ▶ environment variables (if supported by the runtime environment)
- ▶ PDFlib error messages are also generated in EBCDIC format (except in Java).

If you prefer to use input text files in ASCII format you can set the *asciifile* option to *true* (default is *false* on IBM Z and *true* on IBM System i. PDFlib will then expect these files in ASCII encoding. String parameters are still expected in EBCDIC encoding, however.

In contrast, the following items must always be treated in binary mode (i.e., any conversion must be avoided):

- ▶ PDF input and output files
- ▶ TrueType and OpenType font files
- ▶ image files and ICC profiles

3.2 Page Descriptions

3.2.1 Coordinate Systems

PDF's default coordinate system is used within PDFlib. The default coordinate system (or default *user space*) has the origin in the lower left corner of the page, and uses the DTP point as unit:

$$1 \text{ pt} = 1/72 \text{ inch} = 25.4/72 \text{ mm} = 0.3528 \text{ mm}$$

The first coordinate increases to the right, the second coordinate increases upwards. PDFlib client programs may change the default user space by rotating, scaling, translating, or skewing, resulting in new user coordinates. The respective functions for these transformations are *PDF rotate()*, *PDF scale()*, *PDF translate()*, and *PDF skew()*. If the coordinate system has been transformed, all coordinates in graphics and text functions must be supplied according to the new coordinate system. The coordinate system is reset to the default coordinate system at the start of each page.

Using metric coordinates. Metric coordinates can easily be used by scaling the coordinate system. The scaling factor is derived from the definition of the DTP point given above:

```
p.scale(28.3465, 28.3465);
```

After this call PDFlib will interpret all coordinates (except for interactive features, see below) in centimeters since $72/2.54 = 28.3465$.

As a related feature, the *userunit* option in *PDF_begin/end_page_ext()* (PDF 1.6) can be specified to supply a scaling factor for the whole page. Note that user units will only affect final page display in Acrobat, but not any coordinate scaling in PDFlib.

Cookbook A full code sample can be found in the *Cookbook topic* `general/metric_topdown_coordinates`.

Coordinates for interactive elements. PDF always expects coordinates for interactive functions, such as the rectangle coordinates for creating text annotations, links, and file annotations in the default coordinate system, and not in the (possibly transformed) user coordinate system. Since this is very cumbersome PDFlib offers automatic conversion of user coordinates to the format expected by PDF. This automatic conversion is activated by setting the *usercoordinates* option to *true*:

```
p.set_option("usercoordinates=true");
```

Since PDF supports only link and field rectangles with edges parallel to the page edges, the supplied rectangles must be modified when the coordinate system has been transformed by scaling, rotating, translating, or skewing it. In this case PDFlib calculates the smallest enclosing rectangle with edges parallel to the page edges, transform it to default coordinates, and use the resulting values instead of the supplied coordinates.

The overall effect is that you can use the same coordinate systems for both page content and interactive elements when the *usercoordinates* option has been set to *true*.

Visualizing coordinates. In order to assist PDFlib users in working with PDF's coordinate system, the PDFlib distribution contains the PDF file *grid.pdf* which visualizes the coordinates for several common page sizes. Printing the appropriately sized page on transparent material may provide a useful tool for preparing PDFlib development.

You can visualize page coordinates in Acrobat DC as follows:

- ▶ To display cursor coordinates use the following: *View, Show/Hide, Cursor Coordinates*
- ▶ The coordinates will be displayed in the unit which is currently selected in Acrobat. To change the display units proceed as follows: go to *Edit, Preferences, [General...], Units & Guides* and choose one of Points, Inches, Millimeters, Picas, Centimeters.

Note that the coordinates displayed refer to an origin in the top left corner of the page, and not PDF's default origin in the lower left corner. See »Using top-down coordinates«, page 60, for details on selecting a coordinate system which aligns with Acrobat's coordinate display.

Rotating objects. It is important to understand that objects cannot be modified once they have been drawn on the page. Although there are PDFlib functions for rotating, translating, scaling, and skewing the coordinate system, these do not affect existing objects on the page but only subsequently drawn objects.

Rotating text, images, and imported PDF pages can easily be achieved with the *rotate* option of *PDF_fit_textline()*, *PDF_fit_textflow()*, *PDF_fit_image()*, and *PDF_fit_pdi_page()*. Rotating such objects by multiples of 90 degrees inside the respective fitbox can be accomplished with the *orientate* option of these functions. The following example generates some text at an angle of 45° degrees:

```
p.fit_textline("Rotated text", 50.0, 700.0, "rotate=45");
```

Cookbook A full code sample can be found in the *Cookbook* topic `textflow/rotated_text`.

Rotation for vector graphics can be achieved by applying the general coordinate transformation functions *PDF_translate()* and *PDF_rotate()*. The following example creates a rotated rectangle with lower left corner at (200, 100). It translates the coordinate origin to the desired corner of the rectangle, rotates the coordinate system, and places the rectangle at (0, 0). The *save/restore* nesting makes it easy to continue placing objects in the original coordinate system after the rotated rectangle is done:

```
p.save();
  p.translate(200, 100);          /* move origin to corner of rectangle*/
  p.rotate(45.0);               /* rotate coordinates */
  p.rect(0.0, 0.0, 75.0, 25.0); /* draw rotated rectangle */
  p.stroke();
p.restore();
```

Using top-down coordinates. Unlike PDF's bottom-up coordinate system some graphics environments use top-down coordinates which may be preferred by some developers. Such a coordinate system can easily be established using PDFlib's transformation functions. However, since the transformations will also affect text output (text easily appears bottom-up), additional calls are required in order to avoid text being displayed in a mirrored sense.

In order to facilitate the use of top-down coordinates PDFlib supports a special mode in which all relevant coordinates will be interpreted differently. The *topdown* feature has been designed to make it quite natural for PDFlib users to work in a top-down coordinate system. Instead of working with the default PDF coordinate system with the origin (0, 0) at the lower left corner of the page and *y* coordinates increasing upwards, a modified coordinate system will be used which has its origin at the upper left corner of the page with *y* coordinates increasing downwards. This top-down coordinate system for a page can be activated with the *topdown* option of *PDF_begin_page_ext()*:

```
p.begin_page_ext(595.0, 842.0, "topdown");
```

For the sake of completeness we'll list the detailed consequences of establishing a top-down coordinate system below.

»Absolute« coordinates will be interpreted in the user coordinate system without any modification:

- ▶ All method parameters which are designated as »coordinates« in the method descriptions. Some examples: *x, y* in *PDF_moveto()*; *x, y* in *PDF_circle()*, *x, y* (but not *width* and *height!*) in *PDF_rect()*; *llx, lly, urx, ury* in *PDF_create_annotation()*.

»Relative« coordinate values will be modified internally to match the top-down system:

- ▶ Text (with positive font size) will be oriented towards the top of the page;

- ▶ When the manual talks about »lower left« corner of a rectangle, box etc. this will be interpreted as you see it on the page;
- ▶ When a rotation angle is specified the center of the rotation is still the origin (0, 0) of the user coordinate system. The visual result of a clockwise rotation will still be clockwise.

Cookbook A full code sample can be found in the *Cookbook* topic `general/metric_topdown_coordinates`.

3.2.2 Page Size

Cookbook A full code sample can be found in the *Cookbook* topic `pagination/page_sizes`.

Standard page formats. Absolute values and symbolic page size names may be used for the *width* and *height* options in `PDF_begin/end_page_ext()`. The latter are called `<format>.width` and `<format>.height`, where `<format>` is one of the standard page formats (in lowercase, e.g. `a4.width`).

Page size limits. Although PDF and PDFlib don't impose any restrictions on the usable page size, Acrobat implementations suffer from architectural limits regarding the page size. Other PDF interpreters may be able to deal with larger or smaller document formats. The page size limits for Acrobat are shown in Table 3.2. In PDF 1.6 and above the *userunit* option in `PDF_begin/end_page_ext()` can be used to specify a global scaling factor for the page.

Table 3.2 Minimum and maximum page size of Acrobat

PDF viewer	minimum page size	maximum page size
without <i>userunit</i> option (default)	1/24" = 3 pt = 0.106 cm	200" = 14400 pt = 508 cm
with <i>userunit</i> option	3 user units	14 400 user units The maximum value 75 000 for <i>userunit</i> allows page sizes up to 14 400 * 75 000 = 1 080 000 000 points = 381 km

Different page size boxes. While many PDFlib developers only specify the width and height of a page, some advanced applications (especially for prepress work) may want to specify one or more of PDF's additional box entries. PDFlib supports all of PDF's box entries. The following entries, which may be useful in certain environments, can be specified by PDFlib clients (definitions taken from the PDF reference):

- ▶ **MediaBox:** this is used to specify the width and height of a page, and describes what we usually consider the page size.
- ▶ **CropBox:** the region to which the page contents are to be clipped; Acrobat uses this size for screen display and printing.
- ▶ **TrimBox:** the intended dimensions of the finished (possibly cropped) page;
- ▶ **ArtBox:** extent of the page's meaningful content. It is rarely used by application software;
- ▶ **BleedBox:** the region to which the page contents are to be clipped when output in a production environment. It may encompass additional bleed areas to account for inaccuracies in the production process.

PDFlib will not use any of these values apart from recording it in the output file. By default PDFlib generates a MediaBox according to the specified width and height of the page, but does not generate any of the other entries. The following code fragment will start a new page and set the four values of the CropBox:

```
/* start a new page with custom CropBox */
p.begin_page_ext(595, 842, "cropbox={10 10 500 800}");
```

3.2.3 Direct Paths

A path is a shape made of an arbitrary number of straight lines, rectangles, circles, Bézier curves, or elliptical segments. A path may consist of several disconnected sections called subpaths. There are several operations which can be applied to a path:

- ▶ Stroking draws a line along the path, using client-supplied options (e.g., color, line width) for drawing.
- ▶ Filling paints the entire region enclosed by the path, using client-supplied options for filling.
- ▶ Clipping reduces the imageable area for subsequent drawing operations by replacing the current clipping area (which is unlimited by default) with the intersection of the current clipping area and the area enclosed by the path.
- ▶ Merely terminating the path results in an invisible path, which will nevertheless be present in the PDF file. This will only rarely be useful.

Direct Paths. Using the path functions *PDF_moveto()*, *PDF_lineto()*, *PDF_rect()* etc. you can construct a direct path which is written to the current page or another content stream (e.g. a template or Type 3 glyph description). Immediately after constructing the path it must be processed with one of *PDF_stroke()*, *PDF_fill()*, *PDF_clip()* and related functions. These functions consume and delete the path. The only way to use a path multiply is with *PDF_save()* and *PDF_restore()*.

It is an error to construct a direct path without applying any of the above operations to it. PDFlib's scoping system ensures that clients obey to this restriction. If you want to set appearance properties (e.g. color, line width) of a path you must do so before starting any drawing operations. These rules can be summarized as »don't change the appearance within a path description«.

Merely constructing a path doesn't result in anything showing up on the page; you must either fill or stroke the path in order to get visible results:

```
p.set_graphics_option("strokecolor=red");
p.moveto(100, 100);
p.lineto(200, 100);
p.stroke();
```

Most graphics functions make use of the concept of a current point, which can be thought of as the location of the pen used for drawing.

Cookbook A full code sample can be found in the *Cookbook* topic `graphics/starter_graphics`.

3.2.4 Path Objects

Path objects are a convenient and powerful alternative to direct paths. They encapsulate all drawing operations for constructing a path. Unlike direct paths, path objects can be used multiply until they are explicitly destroyed with *PDF_delete_path()*. Informa-

tion about a path can be retrieved with `PDF_info_path()`. Path objects can be created in multiple ways.

Create a path object directly. Calls to `PDF_add_path_point()` add points and associated path elements to a path object. This method can also add a reference to an existing path object to a newly constructed path. `PDF_add_path_point()` supports several convenience options to facilitate path construction. The following code fragment creates a simple path object with a circle, strokes it at two different locations on the page, and finally deletes it:

```
path = p.add_path_point( -1,  0, 100, "move", "");
path = p.add_path_point(path, 200, 100, "control", "");
path = p.add_path_point(path,  0, 100, "circular", "");

p.draw_path(path,  0,  0, "stroke");
p.draw_path(path, 400, 500, "stroke");
p.delete_path(path);
```

Cookbook A full code sample can be found in the *Cookbook* topic `graphics/starter_path`.

Create a path object from an image clipping path. A clipping path contained in a TIFF or JPEG raster image can be retrieved with `PDF_info_image()` and the keyword `clipping-path`:

```
image = p.load_image("auto", "image.tif", "clippingpathname={path 1} infomode");

path = (int) p.info_image(image, "clippingpath", "");
if (path == -1)
    throw new Exception("Error: clipping path not found!");

p.draw_path(path, 0, 0, "stroke");
```

This provides a convenient method for importing arbitrary paths. If you don't need the pixel data of the actual image, you can reduce the output file size by loading the image with the option `infomode`.

Cookbook A full code sample can be found in the *Cookbook* topic `graphics/import_path_from_image`.

Create a path object from SVG path syntax. Path objects can be constructed with the SVG path syntax as required for the `d` attribute of the `path` element (see www.w3.org/TR/SVG11/paths.html#PathData). A full SVG graphics file is not required, but only a String containing the path description which is supplied to the `svgpath` option of `PDF_add_path_point()`. If the path directly stems from an SVG file it uses top-down coordinates which must be adjusted for use in PDFlib (even with the `topdown` page option). First, supply `scale={1 -1}` to unmirror the path. Second, use `PDF_info_path()` to determine the coordinates of that corner of the path's bounding box which will be placed at the target location as shown in the following code fragment where (x_1, y_1) is the lower left corner:

```
String svgPath="M 100 100 L 300 100 L 200 300 z";

/* Create a path object from the SVG path description above */
int path = p.add_path_point(-1, 0, 0, "addpath", "svgpath = {" + svgPath + "} close");
String mirroropt = "scale= {1 -1}";
```

```

double corner_x = p.info_path(path, "x1", mirroropt);
double corner_y = p.info_path(path, "y1", mirroropt);
double target_x = 200, target_y = 500;

/* Place the selected corner at the target location; with the "topdown" page option
   use (target_x-corner_x, target_y+corner_y) instead.
*/
p.draw_path(path, target_x-corner_x, target_y-corner_y,
            mirroropt + " stroke linewidth=3");

```

Cookbook A full code sample can be found in the *Cookbook* topic `graphics/svg_path`.

Create a path object from an object bounding box. The enclosing rectangle (bounding box) of a placed PDF page, SVG graphics, path, raster image, table, matchbox, Textflow or Textline can be retrieved with the corresponding `PDF_info_*`() method:

```

optlist = "boxsize={400 300} fitmethod=clip matchbox={name=border}";
p.fit_image(image, 200, 150, optlist);

int path = (int) p.info_matchbox("border", 1, "boundingbox");
p.draw_path(path, 0, 0, "close stroke linewidth=10 strokecolor=red");
p.delete_path(path);

```

Using path objects. Once a path object has been created or retrieved it can be used for different purposes:

- ▶ `PDF_draw_path()` can be used to fill or stroke the path on the page, or use it as a clipping path.
- ▶ Wrap multi-line Text around the text with `PDF_fit_textflow()`: the text is formatted so that it wraps inside or outside of an arbitrary shape (see Section 9.2.6, »Wrapping Text around Paths and Images«, page 221).
- ▶ Place text on the path with `PDF_fit_textline()`, i.e. the characters follow the curves of the path (see Section 9.1.6, »Text on a Path«, page 203).
- ▶ Place path objects in table cells with `PDF_add_table_cell()`.

3.2.5 Templates (Form XObjects)

Templates (Form XObjects). PDFlib supports a PDF feature with the technical name *Form XObjects*. However, since this term conflicts with interactive forms we refer to this feature as *templates*. A PDFlib template can be thought of as an off-page buffer into which text, vector, and image operations are redirected (instead of directly acting on a page). When the template is finished it can be used like a raster image, and placed an arbitrary number of times on arbitrary pages. Like images, templates can be subjected to geometrical transformations such as scaling or skewing. When a template is used on multiple pages (or multiply on the same page), the actual PDF operators for constructing the template are only included once in the PDF file, thereby saving PDF output file size. Templates are recommended for elements which appear repeatedly on several pages, such as a constant background, a company logo, or graphical elements emitted by CAD and geographical mapping software. Templates are also recommended for raster images with a clipping path if the image is placed more than once. Templates can be created in the following ways:

- ▶ directly with `PDF_begin_template_ext()`;

- ▶ indirectly from vector graphics with `PDF_load_graphics()` and the `templateoptions` option;
- ▶ indirectly from raster images with `PDF_load_image()` and the `templateoptions` option; without this option `PDF_load_image()` creates a similar PDF construct called Image XObject.

Note PDF pages imported with `PDF_open_pdi_page()` also create PDF Form XObjects, but these are handled with PDI functions, not template functions.

Templates can be used in the following ways:

- ▶ place the template on a page or another content stream with `PDF_fit_image()` (see below);
- ▶ create a graphics state with a luminosity soft mask defined by the template (suboption `template` of option `softmask` of `PDF_create_gstate()`, see Section 4.9.2, »Changing the Color with Soft Masks«, page 91);
- ▶ as fallback (background) for loading SVG graphics (option `fallbackimage` of `PDF_load_graphics()`);
- ▶ as appearance of annotations (suboptions `normal/rollover/down` of option `template` of `PDF_create_annotation()`);
- ▶ as appearance of pushbutton form fields (options `icon/icondown/iconrollover` of `PDF_create_field()`)

Creating and using templates with PDFlib. A template can be placed on the page or on another template with the `PDF_fit_image()` method just like raster images (see Section 8.4, »Placing Images, Graphics, and imported PDF Pages«, page 190). The general idiom for creating and using templates in PDFlib looks as follows:

```
/* define the template */
template = p.begin_template_ext(template_width, template_height, "");
...place marks on the template using text, vector, and image functions...
p.end_template_ext(0, 0);
...
p.begin_page(page_width, page_height);
/* use the template */
p.fit_image(template, 0.0, 0.0, "");
...more page marking operations...
p.end_page();
...
p.close_image(template);
```

All text, graphics, and color functions can be used on a template. However, the following functions must not be used while constructing a template:

- ▶ `PDF_begin_item()` and the `tag` option of various functions: structure elements cannot be created within a template.
- ▶ All interactive functions: these must be defined on the page where they should appear in the document, and cannot be generated as part of a template.

Cookbook A full code sample can be found in the *Cookbook* topic `general/repeated_contents`.

3.3 PDF Password Security

3.3.1 Password Security in PDF

PDF password security offers the following protection features:

- ▶ The user password (also referred to as open password) is required to open the file for viewing. Only files with a user password are safe from cracking!
- ▶ The master password (also referred to as owner or permissions password) is required to change any security settings, i.e. permissions, user or master password. Files with user and master passwords can be opened for viewing by supplying either password.
- ▶ Permission settings restrict certain actions for the PDF document, such as printing or extracting text.
- ▶ An attachment password can be specified to encrypt only file attachments, but not the actual contents of the document itself.

If a PDF document uses any of these protection features it is encrypted.

Permission restrictions. PDF can encode various restrictions on document operations which can be granted or denied individually:

- ▶ *Printing Allowed:* If printing is not allowed, the print button in Acrobat remains disabled. Acrobat supports a distinction between *Low Resolution (150 dpi)* and High Resolution printing. Low-resolution printing generates a raster image of the page which is suitable only for personal use, but prevents high-quality reproduction. Note that image-based printing not only results in low output quality, but also considerably slows down the printing process.
- ▶ *Changes Allowed:* the corresponding list provides control over various document modification operations:

Inserting, deleting, and rotating pages

Filling in form fields and signing existing signature fields

Commenting, filling in form fields, and signing existing signature fields

Any except extracting pages

Note that Adobe Reader does not reliably display permission restrictions. For example, since it does not include assembly functions, it always displays *Document Assembly: Not allowed* regardless of the actual permission settings in a document.

- ▶ Content copying is controlled via *Enable copying of text, images, and other content*. While this can be enabled for accessibility with *Enable text access for screen reader devices for the visually impaired*, this setting is deprecated in PDF 2.0 since a PDF reader should always support accessibility.

Specifying access restrictions for a document, such as *Printing allowed: None* disables the respective function in Acrobat. However, this not necessarily holds true for other PDF viewers. It is up to the developer of PDF tools whether or not access permissions are honored.

3.3.2 Password-Protecting PDF Documents with PDFlib

PDFlib can apply standard security features when generating PDF documents. In order to import pages from protected PDF documents with PDFlib+PDI or PDFlib Personalization Server (PPS) the master password or the *shrug* option is required. Querying document properties with the pCOS interface is governed by the pCOS mode. For example,

XMP document metadata, document info fields, bookmarks, and annotation contents can be retrieved without the master password if the document does not require a user password (or only the user password has been supplied). The pCOS Path Reference discusses this in more detail.

Note You cannot Reader-enable PDF documents (e.g. allow annotations with Adobe Reader) with PDFlib GmbH products.

Encryption algorithm and key length. Encryption algorithm and key length used to protect a document with a password depend on the PDF version of the generated document according to the *compatibility* option of `PDF_begin_document()`:

- ▶ PDF 1.4 and 1.5: encryption is not supported since no secure PDF encryption algorithm is available.
- ▶ PDF 1.6, PDF 1.7 and PDF 1.7ext3: AES-128 is used, but encryption for these PDF versions is deprecated. AES-256 according to PDF 1.7ext3 (Acrobat 9) is never used because of its known weaknesses.
- ▶ PDF 1.7ext8 and PDF 2.0: AES-256 according to Acrobat X/XI/DC is used. This is the recommended algorithm. It is used by default since the default PDF output version is PDF 1.7ext8.

Setting passwords with PDFlib. Passwords can be set with the *userpassword* and *masterpassword* options in `PDF_begin_document()`. PDFlib interacts with client-supplied passwords for the generated document in the following ways:

- ▶ If a user password or some permission restrictions, but no master password has been supplied, a regular user would easily be able to change the security settings, thereby defeating any protection. For this reason PDFlib considers this situation as an error.
- ▶ If the user and master password are the same, a distinction between user and owner of the file would no longer be possible, again defeating effective protection. PDFlib considers this situation as an error.
- ▶ Unicode passwords are allowed for AES-256 while AES-128 passwords are restricted to the Latin-1 character set.
- ▶ Passwords are truncated to 127 UTF-8 bytes for AES-256 and 32 characters for AES-128

Setting permissions with PDFlib. Access restrictions can be set with the *permissions* option in `PDF_begin_document()`. It contains one or more access restriction keywords.

When setting the *permissions* option the *masterpassword* option must also be set because otherwise users could easily remove the permission settings. By default, all actions are allowed. Specifying an access restriction disables the respective feature in Acrobat. Access restrictions can be applied without any user password. Multiple restriction keywords can be specified as in the following example:

```
p.begin_document(filename, "masterpassword=abcd1234 permissions={noprint nocopy}");
```

Table 3.3 lists all supported access restriction keywords.

Cookbook A full code sample can be found in the Cookbook topic `general/permission_settings`.

Table 3.3 Access restriction keywords for the *permissions* option of `PDF_begin_document()`

keyword	explanation
<code>noprint</code>	Acrobat prevents printing the file.

Table 3.3 Access restriction keywords for the permissions option of `PDF_begin_document()`

keyword	explanation
<code>nomodify</code>	Acrobat prevents users from adding form fields or making any other changes.
<code>nocopy</code>	Acrobat prevents copying and extracting text or graphics and disables accessibility.
<code>noannots</code>	Acrobat prevents adding or changing comments or form fields.
<code>noforms</code>	(Implies <code>noannots</code>) Acrobat prevents form field filling, even if <code>noannots</code> hasn't been specified.
<code>noaccessible</code>	(Deprecated and not available in PDF 2.0) Acrobat prevents extracting text or graphics for accessibility purposes.
<code>noassemble</code>	(Implies <code>nomodify</code>) Acrobat prevents inserting, deleting, or rotating pages and creating bookmarks and thumbnails, even if <code>nomodify</code> hasn't been specified.
<code>nohighresprint</code>	Acrobat prevents high-resolution printing. If <code>nohighresprint</code> hasn't been specified printing is restricted to the »print as image« feature which prints a low-resolution rendition of the page.
<code>plainmetadata</code>	Keep document metadata unencrypted even for encrypted documents.

Encrypted file attachments. File attachments can be encrypted even in otherwise unprotected documents. This can be achieved by supplying the `attachmentpassword` option to `PDF_begin_document()`.

3.4 XMP Metadata

As an alternative or in addition to document information fields PDFlib supports XMP (*Extensible Metadata Platform*) as a framework for specifying metadata. XMP has been standardized as ISO 16684-1:2012. There are several aspects of XMP support in PDFlib as detailed below.

Cookbook A simple XMP sample can be found in the *Cookbook* topic `interchange/embed_xmp`.

Most commonly XMP is used to attach metadata to the whole document. In addition to document-level metadata, XMP can be supplied for pages, fonts, images, graphics, templates, and imported PDF pages. This can be achieved with the `metadata` option of various methods, for example:

```
metadata={filename=info.xmp inputencoding=winansi}
```

The `metadata` option expects full or partial XMP metadata streams. PDFlib validates the user-supplied XMP metadata according to XML and XMP/RDF rules. For PDF/A additional rules for custom XMP properties apply; see Section 12.2.8, »XMP Metadata for PDF/A«, page 314.

Internal and reserved XMP properties. PDFlib creates several XMP properties internally, e.g. `CreationDate`. Other XMP properties are required to signal conformance to various PDF standards such as PDF/A or PDF/X. Internal properties as well as standards-related identification properties can not be overridden by user-supplied XMP.

Automatic XMP synchronization for document info fields. If the `autoxmp` option in `PDF_begin/end_document()` is `true`, PDFlib synchronizes document information fields

supplied to `PDF_set_info()` as well as several internally generated entries (e.g. *Creation-Date*) to the corresponding entries in the document-level XMP metadata.

Document info fields which correspond to a well-known property in one of the standard XMP schemas are placed in the appropriate schema. Custom info fields are placed in the extended PDF (*pdfx*, unrelated to PDF/X) XMP schema, but are ignored in PDF/A mode since custom XMP properties require an extension schema description in PDF/A.

Cloning XMP metadata. If most or all pages of a PDF document are imported it is recommended to clone XMP metadata if it is present in the input. XMP metadata can be cloned with the following code fragment:

```
if (p.pcos_get_string(indoc, "type:/Root/Metadata").equals("stream"))
{
    xmp = p.pcos_get_stream(indoc, "", "/Root/Metadata");
    p.create_pvf("/xmp/document.xmp", xmp, "");
    optlist += " metadata={filename=/xmp/document.xmp}";
}

p.end_document(optlist);
p.delete_pvf("/xmp/document.xmp");
```

3.5 Associated Files

Note More information about associated files can be found in Application Note 0002, published by the PDF Association (www.pdfa.org/resource/pdf-2-0-application-note-002-associated-files/)

Associated files are similar to file attachments, but can be used to connect data with individual objects in a PDF document. Associated files allow you to relate content in arbitrary formats to PDF objects and to specify the type of relationship. One or more files can be associated with a particular object. This can be achieved with the option *associatedfiles* of various API methods.

Associated files must be loaded with `PDF_load_asset()`. In contrast to regular file attachments, associated files are by default not listed in Acrobat's attachments pane. If you load the asset with the option *documentattachment* the associated file is included in the list of document attachments. As a result, the associated file(s) are visible in the list of attachments. The files are embedded in the generated PDF document unless the option *external=true* is supplied.

Associated files are part of PDF 2.0, but have been introduced with PDF/A-3 which imposes the following restrictions on associated files:

- ▶ references to unembedded external files, i.e. *external=true* are not allowed;
- ▶ associated files for the structure hierarchy (document option *structureassociatedfiles*) and XMP metadata are not allowed.

Table 3.4 lists all API methods and options with support for associated files. The use cases also mention typical values of the *relationship* option.

Table 3.4 API methods with support for the associatedfiles option or suboption

API method	file is associated with...	sample use cases
<i>PDF_end_document()</i>	whole document	ZUGFeRD and Factur-X embed an XML version of the invoice into a PDF/A-3 document (relationship=Alternative); word processor or spreadsheet source file for the document (relationship=Source)
<i>PDF_begin/end_page_ext()</i>	page	associated files may serve as an alternative to a vendor-specific PieceInfo dictionary
<i>PDF_begin/end_document()</i> , <i>PDF_begin/end_page_ext()</i> and other API methods with option metadata	XMP metadata	schema description for XMP metadata (relationship=Schema)
<i>PDF_load_graphics()</i> with <i>templateoptions</i> <i>PDF_open_pdi_page()</i> <i>PDF_begin_template_ext()</i>	FormXObject created for the template or the imported page or graphics	source data for a chart, CSV data for a table, raw measurements for a graph (relationship=Data); MathML representation of a formula (relationship=Alternative)
<i>PDF_load_image()</i>	ImageXObject created for the imported image	source data for an image (relationship=Source)
<i>PDF_create_annotation()</i>	annotation (with type=FileAttachment and option attachment the embedded file is associated with the page instead)	source data for RichMedia or 3D annotations (relationship=Source)
<i>PDF_begin_item()</i> or tag option of various methods	structure element in the logical structure hierarchy	Responsive PDF embeds HTML, SVG, JavaScript or CSS for styling the structure element contents when converting to another format (relationship=Alternative or Supplement)
<i>PDF_end_document()</i> with <i>structureassociatedfiles</i>	entire logical structure hierarchy	supplementary information for converting the structure tree to another format (relationship=Supplement)
<i>PDF_begin/end_dpart()</i>	document part in the DPart hierarchy	additional information about the document part (relationship=Supplement)

Cookbook A full code sample for connecting associated files with various objects within a document can be found in the topic *interchange/associated_files* in the PDFlib Cookbook. A code sample for retrieving associated files from a PDF document can be found in the topic *interchange/retrieve_associated_files* in the pCOS Cookbook.

The following code fragment loads a file asset and associates it with a page while specifying the relationship as *Data*:

```
af = p.load_asset("Attachment", af_filename, "relationship=Data");

if (af == -1) {
    throw new Exception("Error: " + p.get_errmsg());
}
```

```
p.begin_page_ext(0, 0, "associatedfiles={" + af + "} width=a4.width height=a4.height");
```



4 Color and Transparency

Color spaces can be specified separately for filling and stroking operations, e.g. with the *fillcolor* and *strokecolor* options. While color spaces in PDF are also used for images, PDFlib usually determines the image color space automatically unless color-related options are supplied in *PDF_load_image()*.

Cookbook Code samples for color handling can be found in the color category of the PDFlib Cookbook. The color/starter_color sample demonstrates the use of all color spaces.

4.1 Device Color Spaces

Device-specific color spaces are perhaps the most common methods for describing color. As the name implies they describe color as rendered by a specific device such as a monitor, printer or printing machine. Device-specific colors appear differently depending on which device produces the output. Because of this undesirable property device colors are not allowed directly in PDF/A and PDF/X, but only with a suitable output intent ICC profile or default color space.

Grayscale color space. Grayscale colors can be requested with the color space keyword *gray* and a single gray value. Gray values are in the range 0=black and 1=white. Example:

```
p.set_graphics_option("fillcolor={{ gray 0.5 }}");
```

RGB color space. Colors in the additive RGB color space are mixed from the red, green and blue components. RGB colors can be requested with the color space keyword *rgb* and three RGB values in the range 0..1 specifying the percentage of red, green, and blue with (0, 0, 0)=black and (1, 1, 1)=white. The commonly used RGB color values in the range 0–255 must be divided by 255 in order to scale them to the range 0–1 as required by PDFlib.

As an alternative to numerical RGB values you can specify RGB colors via their HTML names or hexadecimal values. Examples:

```
p.set_graphics_option("fillcolor={{ rgb 1 0 0 }}");  
p.set_graphics_option("fillcolor=pink");  
p.set_graphics_option("fillcolor={#FFC0CB}");
```

CMYK color space. Colors in the subtractive CMYK color space are mixed from the cyan, magenta, yellow and black («key») components. CMYK values can be requested with the color space keyword *cmymk* and four CMYK values in the range 0..1, where 0 = no color and 1 = full color, representing cyan, magenta, yellow, and black values; (0, 0, 0, 0)=white, (0, 0, 0, 1)=black. Note that the polarity of CMYK colors is different from grayscale and RGB colors. Example:

```
p.set_graphics_option("fillcolor={{ cmyk 0 1 0 0 }}");
```

A special property of the CMYK color space is that the overprint behavior of CMYK objects can be controlled with the *overprintmode* option (see «Overprint mode for CMYK colors», page 93).

Mapping device colors to device-independent colors with default color spaces. Device color spaces can be mapped to device-independent color spaces by means of default color spaces; see »Mapping device colors to ICC-based color spaces«, page 76, for details. This is useful when data with device-specific colors, such as RGB or CMYK images or graphics, is used in PDF/A or PDF/X without any matching output intent.

4.2 Color Management with ICC Profiles

The International Color Consortium (ICC) defined a file format for specifying color characteristics of input and output devices. PDFlib supports color management with ICC profiles and rendering intents. ICC profiles play a crucial role in color-managed workflows and most PDF standards such as PDF/X and PDF/A.

Cookbook Full code samples can be found in the `starter_color` sample and the *Cookbook* topic `color/iccprofile_to_image`.

4.2.1 Loading ICC Profiles

Note Recommendations and links to freely available ICC color profiles for common printing conditions are available at www.pdflib.com.

Acceptable ICC profiles. ICC profiles can be loaded with `PDF_load_iccprofile()`. Profiles must satisfy certain conditions regarding the ICC version number of the profile, device class and data color space. PDF 1.4 supports ICC profile versions 2.x; PDF 1.5 and above supports ICC profile versions 4.x. Table 4.1 details additional requirements regarding device class and data color space for ICC profiles depending on their intended use.

Table 4.1 ICC profile requirements for various uses

ICC profile usage (see Table 4.2 for details)	device class	data color space
output intent for PDF/X-3/4	prtr	Gray, RGB, CMYK
output intent for PDF/X-5n	prtr	xCLR (n-colorant)
document or page-level output intent for PDF 2.0 and PDF/A	prtr, mntr	Gray, RGB, CMYK
transparency group color space	prtr, mntr, scnr, spac	Gray, RGB, CMYK
all other uses of ICC profiles	prtr, mntr, scnr, spac	Gray, RGB, CMYK

Searching ICC profiles. PDFlib searches for ICC profiles according to the following steps, using the `profilename` parameter supplied to `PDF_load_iccprofile()`:

- ▶ If `profilename=sRGB`, PDFlib uses the internal sRGB profile (see below) and terminates the search.
- ▶ Check whether there is a resource named `profilename` in the `ICCProfile` resource category. If so, use its value as file name in the following steps. If there is no such resource, use `profilename` as a file name directly.
- ▶ Use the file name determined in the previous step to locate a disk file by trying the following combinations one after another:

```
<filename>  
<filename>.icc
```

```

<filename>.icm
<colordir>/<filename>(only on Windows and macOS)
<colordir>/<filename>.icc(only on Windows and macOS)
<colordir>/<filename>.icm(only on Windows and macOS)

```

On Windows *colordir* designates the directory where device-specific ICC profiles are stored by the operating system (e.g. *C:\Windows\System32\spool\drivers*). On macOS the following paths are tried for *colordir*:

```

/System/Library/ColorSync/Profiles
/Library/ColorSync/Profiles
/Network/Library/ColorSync/Profiles
~/Library/ColorSync/Profiles

```

The sRGB color space and sRGB ICC profile. PDFlib supports a standard RGB color space called sRGB. It is widely used for simplified color management for consumer RGB devices such as digital cameras as well as office equipment such as color printers and monitors. PDFlib includes the required ICC profile data internally. Therefore the sRGB profile must not be configured explicitly. It can be requested by calling *PDF_load_iccprofile()* with *profilename=sRGB*. As a convenient shortcut the keyword *srgb* can be supplied as an alternative in all places where an ICC profile handle created with *PDF_load_iccprofile()* is expected. The sRGB profile is applied to RGB images by default unless the image file contains an embedded ICC profile.

The sRGB profile is of device class *mnr* (output device), i.e. it can be used as output intent for PDF/A, but not for PDF/X.

4.2.2 Using ICC Profiles

PDFlib supports ICC profiles for the use cases listed in Table 4.2.

Table 4.2 Different uses of ICC profiles

Use case	Relevant API methods and options
Set ICC-based color space for text and vector graphics on the page	<i>PDF_set_option()</i> with <i>iccprofilegray/rgb/cmyk</i> and <i>PDF_setcolor()</i> with <i>colorspace=iccbasedgray/rgb/cmyk</i> ; color options with <i>iccbased</i> keyword
Apply ICC profile to an imported image	<i>PDF_load_image()</i> : option <i>iccprofile</i>
Process or ignore ICC profile embedded in an image	<i>PDF_load_image()</i> : option <i>honoriccprofile</i>
Query ICC profile embedded in an image	<i>PDF_info_image()</i> with <i>keyword=iccprofile</i>
Set default color space for mapping grayscale, RGB, or CMYK data to ICC-based color spaces	<i>PDF_begin_page_ext()</i> , <i>PDF_begin_template_ext()</i> , <i>PDF_begin_pattern_ext()</i> and <i>PDF_begin_font()</i> : options <i>defaultgray/defaultrgb/defaultcmyk</i>
Specify a PDF/X or PDF/A document output intent with a referenced or embedded ICC profile	<i>PDF_load_iccprofile()</i> : option <i>usage=outputintent</i>
Specify a PDF 2.0 page-level output intent with an embedded ICC profile	<i>PDF_begin_page_ext()</i> : option <i>outputintents</i>
Specify a blending color space for transparency and blend modes	<i>PDF_begin/end_page_ext()</i> , <i>PDF_open_pdi_page()</i> , <i>PDF_begin_template_ext()</i> and <i>PDF_load_graphics()</i> with <i>templateoptions: option transparencygroup, suboption colorspace</i>

Table 4.2 Different uses of ICC profiles

Use case	Relevant API methods and options
Alternate color space for spot color	<code>PDF_set_graphics_option()</code> : color option with keyword <code>spotname</code> and icc-based <code>alternate_color</code> , or <code>PDF_set_graphics_option()</code> with option <code>fillcolor</code> followed by <code>PDF_makespotcolor()</code> .
Alternate color space for DeviceN color	<code>PDF_create_devicen()</code> : option <code>alternate</code> with suboption <code>iccbased</code>
Query number of color components in an ICC profile	<code>PDF_get_option()</code> with keyword <code>icccomponents</code>

Using embedded profiles in images (ICC-tagged images). Some images contain embedded ICC profiles describing the nature of the image's color values. For example, an embedded ICC profile can describe the color characteristics of the scanner used to produce the image data. PDFlib handles embedded ICC profiles in the JPEG, JPEG 2000, PNG and TIFF image file formats. If the `honoriccprofile` option is set to `true` (which is the default) the ICC profile embedded in an image is extracted from the image and embedded in the PDF output.

The keyword `iccprofile` of `PDF_info_image()` can be used to obtain an ICC profile handle for the profile embedded in an image. This may be useful when the same profile must be applied to multiple images.

In order to check the number of color components in an unknown ICC profile use the `icccomponents` option.

Applying external ICC profiles to images. As an alternative to using ICC profiles embedded in an image, an external profile may be applied to an individual image by supplying a profile handle along with the `iccprofile` option to `PDF_load_image()`.

ICC-based color spaces for page descriptions. The color values for text and vector graphics can directly be specified in an ICC-based color space. In color options you specify the ICC profile by handle, followed by the appropriate number of color components (e.g. four components for a CMYK profile):

```
icchandle = p.load_iccprofile("myCMYK", "usage=iccbased errorpolicy=exception");
p.set_graphics_option("fillcolor={iccbased=" + icchandle + " 0 1 0 0}");
```

Alternatively, the color space can be set by supplying the ICC profile handle as value to one of the global `iccprofilegray`, `iccprofilergb`, `iccprofilecmyk` options. Subsequently ICC-based color values can be supplied to a color option or `PDF_setcolor()` along with one of the color space keywords `iccbasedgray`, `iccbasedrgb`, or `iccbasedcmyk`.

Mapping device colors to ICC-based color spaces. PDF provides a feature for mapping device-dependent gray, RGB, or CMYK colors to device-independent ICC-based colors. This can be used to attach a precise colorimetric specification to color values which otherwise would be device-dependent. It can be achieved by supplying the `defaultgray`, `defaultrgb`, or `defaultcmyk` options of `PDF_begin_page_ext()`, `PDF_begin_template_ext()`, `PDF_begin_pattern_ext()` and `PDF_begin_font()` with a suitable ICC profile handle. The following example sets the sRGB color space as the default RGB color space for text, images, and vector graphics on the page:

```
p.begin_page_ext(595, 842, "defaultrgb=srgb");
```

If the default color space is derived from an external ICC profile, a profile handle must first be created:

```
/* create ICC profile handle*/  
icchandle = p.load_iccprofile("myRGB", "usage=iccbased errorpolicy=exception");  
p.begin_page_ext(595, 842, "defaultrgb=" + icchandle);
```

Output intent profiles . An output device profile can be used to specify an output condition for a PDF 2.0, PDF/X or PDF/A document. This is done by supplying *usage=output-intent* in `PDF_load_iccprofile()`. For PDF/A a printer or monitor profile can be specified as output intent, while PDF/X allows only printer profiles. For details see Section 12.3, »PDF/X for Print Production«, page 316, and Section 12.2, »PDF/A for Archiving«, page 305.

In PDF 2.0 output intent profiles can also be specified for individual pages with *usage=pageoutputintent* in `PDF_load_iccprofile()` and *outputintents* in `PDF_begin_page_ext()`.

4.2.3 Advanced Color Management Controls

Although PDFlib applications can specify device-independent color values, a particular output device may not be able to accurately reproduce the colors. In this situation some compromises have to be made in a process called gamut compression, i.e., reducing the range of colors to a smaller range which can be reproduced by the device. PDF offers several controls to fine-tune this process.

Rendering intents. Rendering intents can be specified for individual images by supplying the *renderingintent* option to `PDF_load_image()`. Rendering intents can also be specified for text and vector graphics by supplying the *renderingintent* option to `PDF_create_gstate()`. Table 4.7 lists all available rendering intents.

Table 4.3 Available rendering intents

<i>intent</i>	<i>description</i>	<i>typical use</i>
Auto	<i>Do not specify any rendering intent in the PDF file, but use the device's default intent instead.</i>	<i>not enough information about the page contents is available</i>
AbsoluteColorimetric	<i>No correction is made for the device's white point (such as paper white). Colors which are out of gamut are mapped to nearest value within the device's gamut.</i>	<i>exact reproduction of solid colors; not recommended for other uses</i>
RelativeColorimetric	<i>The color data is compressed into the device's gamut, mapping the white points onto one another (i.e. the lightest color is usually produced by paper white). The remaining colors are slightly shifted.</i>	<i>vector graphics</i>
Saturation	<i>Color saturation is preserved while the color values may be shifted.</i>	<i>business graphics</i>
PerceptualColor	<i>Color relationships are preserved by modifying both in-gamut and out-of-gamut colors in order to provide a pleasing appearance.</i>	<i>real-world images</i>

Black point compensation. Black point compensation adjusts the color transformation such that darkest colors are mapped to the darkest color that the output device can produce. This is a similar effect on black as what the *RelativeColorimetric* rendering intent does with white: it maps the lightest color to the lightest color that the output device can produce.

Black point compensation can enhance image quality by achieving a higher contrast. It is available only in PDF 2.0 and can only be enabled with the *useblackptcomp* option of *PDF_create_gstate()*.

4.3 Device-Independent CIE L*a*b* Color

Cookbook A full code sample can be found in the `starter_color` sample.

Device-independent color values can be specified in the CIE 1976 L*a*b* color space (*Lab* for short). Colors in the L*a*b* color space are specified by three values *L*, *a* and *b* (see figure 4.1). The luminance (or lightness) *L* runs in the range 0-100. The values *a* and *b* in the range -128 to 127 describe the color. The value *a* runs from green (-128) to magenta-red (+127); *b* runs from blue (-128) to yellow (+127). Positive numbers describe warm colors (yellow, magenta-red), while negative numbers describe cool colors (green, blue). The intersection of the *a* and *b* axes (*a=b=0*) describes neutral gray values ranging from black (0, 0, 0) to white (100, 0, 0).

An important property of the Lab color space is that unlike RGB and CMYK it is perceptually uniform, which means that the same numerical difference corresponds to the same visual difference between colors. This makes it much easier to calculate mixtures of multiple colors simply by calculating the weighted average of their Lab values. We will take advantage of this when creating a PostScript transform function in »DeviceN color space based on spot colors with Lab alternate space«, page 84.

Lab values specify absolute colors only if the white point is also specified. PDFlib uses the white point of the standard illuminant D50 (daylight 5000 K, 2° observer).

PDFlib supports the Lab color space as follows:

- ▶ Lab colors for text and vector graphics can be specified in color options and *PDF_setcolor()* using the color space keyword *lab* plus three numbers for the *L*, *a* and *b* values, e.g.

```
p.set_graphics_option("fillcolor={lab 100 0 0}");
```

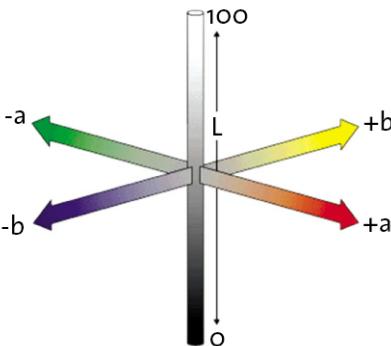


Fig. 4.1
The CIE L*a*b* color space

In PDF/A mode the default black fill and stroke colors are specified as device-independent Lab (0, 0, 0) unless an output intent has been specified.

- ▶ PDFlib uses the Lab color space for the alternate color values of all HKS and Pantone spot colors in its internal database.
- ▶ TIFF images with Lab color space can be imported.
- ▶ ICC profiles with Lab color space can be loaded with `PDF_load_iccprofile()` unless they are intended for use as output intent or for the `transparencygroup` option since PDF doesn't allow Lab in these places.

4.4 Pantone, HKS, and custom Spot Colors

Cookbook Full code samples can be found in the `starter_color` sample and the *Cookbook* topic `color/spot_color`.

Separation color space for spot colors. PDFlib supports spot colors (technically known as Separation color space in PDF) which can be used to print custom colors outside the range of colors mixed from process colors. Spot colors are specified by name, and in PDF must be accompanied by an alternate color which closely, but not exactly, resembles the spot color. The PDF viewer or RIP uses the alternate color for screen display and printing to devices which do not support spot colors (such as office printers). On the printing press the requested spot color is applied in addition to any process colors which may be used in the document.

PDFlib supports various built-in spot color libraries as well as custom (user-defined) spot colors. When a spot color name is requested with `PDF_makespotcolor()` PDFlib first checks whether the requested spot color can be found in one of its built-in libraries. If so, PDFlib uses built-in Lab values for the alternate color. Otherwise the spot color is assumed to be a user-defined color, and the client must supply appropriate alternate color values in the `spotname` color keyword of color options (or via the current fill color when using `PDF_setcolor()`). Spot colors can be tinted, i.e., they can be used with a percentage which specifies the color intensity from 0=no color=white to 1=maximum intensity. Separation color spaces are always subtractive (unlike additive color spaces like RGB where 0=black). Once a spot color has been defined, it can be used to draw text and vector graphics, to colorize grayscale images, or to construct a shading. Note that PDFlib automatically inverts the color polarity when colorizing images with a spot color (see »Colorize a grayscale image with a spot color«, page 168).

By default, built-in spot colors cannot be redefined with custom alternate values. This behavior can be changed with the `spotcolorlookup` option of `PDF_set_option()`. This can be useful to achieve compatibility with older applications which may use different color definitions, and for workflows which cannot deal with PDFlib's Lab alternate values for Pantone colors.

Note Built-in Pantone® and HKS® spot color data and the corresponding trademarks have been licensed by PDFlib GmbH from the respective trademark owners for use in PDFlib software.

Pantone® colors. Pantone colors are well-known and widely used on a world-wide basis. PDFlib fully supports the Pantone Matching System® with thousands of named colors. All color swatch names from the digital color libraries listed in Table 4.4 can be used.

Spot color names are case-sensitive; use uppercase as shown in the examples. The *PANTONE* prefix must always be provided in the swatch name as shown in the examples. Pantone color names are constructed according to the following scheme:



```
PANTONE <id> <paperstock>
```

where *<id>* is the identifier of the color (e.g., 185) and *<paperstock>* the abbreviation of the paper stock or another designation (e.g., C for coated). A single space character must be provided between all components constituting the swatch name. If a spot color is requested where the name starts with the *PANTONE* prefix, but the name does not represent a valid Pantone color a warning is logged. The following code snippet demonstrates the use of a Pantone color with a tint value of 70 percent:

```
p.set_graphics_option("fillcolor={ spotname {PANTONE 281 U} 0.7 }");
```

Table 4.4 Pantone spot color libraries built into PDFlib

color library name	sample color name	remarks
<i>PANTONE solid coated</i>	<i>PANTONE 185 C</i>	336 colors in the range 2001-2336 added in 2012; 84 colors in the range 2337-2427 added in 2014; 112 colors in the range 2428-3599 added in 2016; 294 colors in the range 4001-4294 added in 2019;
<i>PANTONE solid uncoated</i>	<i>PANTONE 185 U</i>	336 colors in the range 2001-2336 added in 2012; 84 colors in the range 2337-2427 added in 2014; 112 colors in the range 2428-3599 added in 2016; 294 colors in the range 4001-4294 added in 2019;
<i>PANTONE Extended Gamut Coated (XGC)</i>	<i>PANTONE 185 XGC</i>	1736 colors introduced in 2015
<i>PANTONE pastel coated</i>	<i>PANTONE 9461 C</i>	
<i>PANTONE pastel uncoated</i>	<i>PANTONE 9461 U</i>	
<i>PANTONE metallic coated</i>	<i>PANTONE 871 C</i>	

HKS® colors. The HKS color system is widely used in Germany and other European countries. PDFlib fully supports HKS colors. All color swatch names from the digital color libraries listed in Table 4.5 can be used.

Spot color names are case-sensitive; use uppercase as shown in the examples. The HKS prefix must always be provided in the swatch name as shown in the examples. HKS color names are constructed according to the following scheme:



```
HKS <id> <paperstock>
```

where `<id>` is the identifier of the color (e.g., 43) and `<paperstock>` the abbreviation of the paper stock in use (e.g., *N* for natural paper). A single space character must be provided between the *HKS*, `<id>`, and `<paperstock>` components constituting the swatch name. If a spot color is requested where the name starts with the *HKS* prefix, but the name does not represent a valid *HKS* color a warning is logged. The following code snippet demonstrates the use of an *HKS* color with a tint value of 70 percent:

```
p.set_graphics_option("fillcolor={ spotname {HKS 38 E} 0.7 }");
```

Table 4.5 *HKS spot color libraries built into PDFlib*

color library name	sample color name	remarks
<i>HKS K</i>	<i>HKS 43 K</i>	88 colors for gloss art paper (<i>Kunstdruckpapier</i>)
<i>HKS N</i>	<i>HKS 43 N</i>	86 colors for natural paper (<i>Naturpapier</i>)
<i>HKS E</i>	<i>HKS 43 E</i>	88 colors for continuous stationary/coated (<i>Endlospapier</i>)
<i>HKS Z</i>	<i>HKS 43 Z</i>	50 colors for newsprint (<i>Zeitungspapier</i>)

Custom spot colors. In addition to built-in spot colors as detailed above, PDFlib supports custom spot colors. These can have an arbitrary name (which must not conflict with the name of any built-in color, however) and an alternate color which will be used for screen preview or low-quality printing, but not for high-quality color separations. The client is responsible for providing suitable alternate colors for custom spot colors.

Spot colors can be set with the *fillcolor/strokecolor* text or graphics appearance options and other color-related options. The alternate color can be supplied directly in the spot color definition:

```
fillcolor={spotname {CompanyRed} 1.0 {cmyk 0 0.78 0.88 0}}
```

Alternatively, spot colors can be defined by using the current fill color as alternate color. Except for an additional call to set the current fill color as alternate color, defining and using custom spot colors works similarly to using built-in spot colors:

```
/* set current fill color for use as alternate color */
p.setcolor("fill", "cmyk", 0 0.78 0.88 0);
/* derive spot color from the current fill color */
spot = p.makespotcolor("CompanyRed");

/* set spot color for filling */
p.set_graphics_option("fillcolor={spotname {CompanyRed} 0.7 }");
```

Separation color space based on CMYK process colors. While spot color names generally can be chosen arbitrarily, the colorant names *Cyan*, *Magenta*, *Yellow* and *Black* always refer to the CMYK process colors. This fact can be used to paint into a specific CMYK channel only (e.g. *Magenta*), without affecting the other channels. This differs from painting in full CMYK color with *Cyan=Yellow=Black=0* which erases existing contents in the *Cyan*, *Yellow* and *Black* channels (this behavior can be modified with device-dependent settings, see Section 4.10, »Overprint Control«, page 93). For example, the following call:

```
fillcolor={spotname Magenta 0.5 {cmyk 0 1 0 0}}
```

is different from directly using CMYK color:

```
fillcolor={cmyk 0 0.5 0 0}
```

since the former call leaves existing contents in the Cyan, Yellow and Black channels unchanged, while the second calls sets them to 0% (subject to the overprint settings).

4.5 DeviceN Colors

Cookbook Full code samples can be found in the `starter_color` sample and the *Cookbook* topic `color/devicen_color`.

DeviceN color spaces support an arbitrary number of named color components and can be regarded as a generalization of spot colors. The colorants may be taken from the set of process colors (typically CMYK, but other process color spaces are possible) or may be arbitrary spot colors. Applications of DeviceN color include the following:

- ▶ Some printing systems use more than four process colors to extend the color gamut (the set of printable colors). For example, some seven-color printing systems are based on CMYK plus Orange, Green and Violet.
- ▶ A DeviceN color space containing a subset of the CMYK process colors, e.g. only Cyan and Magenta, can be used when an object in CMYK color should overprint other CMYK objects.
- ▶ DeviceN color can be used to construct a shading (smooth color transition) between a spot color and a process color or between multiple spot colors.
- ▶ Package printing often uses n-colorant color because traditional spot colors cannot be used. Additional color channels are also used to record information which is not directly related to visible colors, such as varnish or die lines.

The method `PDF_create_devicen()` returns a DeviceN color space handle which can be used for drawing operations with color-related options such as `fillcolor/strokecolor`, for constructing a shading with `PDF_shading()`, or for coloring a raster image with the `colorize` option of `PDF_load_image()`. DeviceN colors can be tinted, i.e., they can be used with a percentage for each color channel which specifies the intensity of the respective color channel from 0=no color=white to 1=maximum intensity. DeviceN color spaces are always subtractive (unlike additive color spaces like RGB where 0=black).

`PDF_create_devicen()` needs the list of colorant names, the alternate color space, and PostScript code which implements the tint transform function. The transform function must convert N color values of the DeviceN color space to the corresponding color values in the alternate color space. The alternate color values are used for rendering if the output device doesn't support the named colors in the DeviceN color space, e.g. on a monitor. Below we provide suitable PostScript transform functions for DeviceN color spaces constructed from CMYK process colors or from spot colors when using Lab as alternate color space. Devising a suitable PostScript transform function for other combinations is a non-trivial task.

When creating output for PDF 1.6 or above, PDFlib emits a DeviceN attributes dictionary with *Colorants* entries for all known spot colors. This ensures that the PDF viewer has enough information to render the named colorants individually and can blend the colorants. The advantage of this method is that the color display in Acrobat is correct regardless of the Overprint Preview setting. It also means that Acrobat doesn't need the PostScript transform function for rendering the page. However, the transform function

must nevertheless be provided since third-party viewers or RIPs may resort to the PostScript function for rendering the DeviceN color to the alternate color space.

In PDF/X-4/5 mode `PDF_makespotcolor()` must be called before `PDF_create_devicen()` for all custom spot colors used in the DeviceN color space. Since this generally improves the screen display it is recommended even if no PDF/X-4/5 is generated.

NChannel color spaces. PDF 1.6 introduced an extension of DeviceN color spaces called NChannel. NChannel color spaces contain additional information which aids the PDF viewer or RIP in accurate color rendering. This information includes alternate color spaces for individual colorants in the DeviceN color space (not only the combined tint transform function), the process color space and component names. NChannel may also include additional information describing the mixing behavior of printing inks, but this is not currently supported in PDFlib. NChannel color spaces require the *Colorants* entry with individual alternate colors for the colorants. This is created automatically by PDFlib. NChannel color spaces also require the name of the process color space and its colorant names. These must be supplied in the *process* option and the *colorspace* and *components* suboptions.

NChannel color spaces can be created with the option `subtype=nchannel` in `PDF_create_devicen()`. The advantage of this setting is that DeviceN color spaces with this subtype no longer depend on Acrobat's *Overprint Preview* setting. In some cases this Acrobat preference is required for correct screen preview, but NChannel automatically forces Acrobat into this display mode.

DeviceN color space based on CMYK process colors. In the following example we use a subset of CMYK colors consisting of the Magenta and Yellow colorants, and DeviceCMYK as alternate color space. The transform function is simple as it only supplies additional zero values for the Cyan and Black components of the CMYK alternate color space:

```
devicen = p.create_devicen(
    "names={Magenta Yellow} alternate=devicecmymk transform={{0 0 4 1 roll}}");
p.set_graphics_option("fillcolor={devicen " + devicen + " 0.5 1}");
```

N color values must be supplied when setting the color (*N*=2 in the example above). The DeviceN color space created above paints only in the Magenta and Yellow channels, but leaves the Cyan and Black channels unchanged. This differs from painting in full CMYK color with *Cyan=Black=0* which erases existing contents in the Cyan and Black channels (subject to the overprint settings).

For convenience Table 4.4 lists the DeviceN PostScript transform functions for all possible subsets of CMYK colorants with CMYK as alternate color space.

Table 4.6 DeviceN PostScript transform functions for subsets of CMYK process colorants and alternate=devicecmymk

DeviceN colorant names for CMYK subsets	PostScript transform function	DeviceN colorant names for CMYK subsets	PostScript transform function
Cyan	{0 0 0}	Magenta Black	{0 3 1 roll 0 exch}
Magenta	{0 0 0 4 1 roll}	Yellow Black	{0 0 4 2 roll}
Yellow	{0 0 0 4 2 roll}	Cyan Magenta Yellow	{0}
Black	{0 0 0 4 3 roll}	Cyan Magenta Black	{0 exch}

Table 4.6 DeviceN PostScript transform functions for subsets of CMYK process colorants and alternate=devicecmymk

DeviceN colorant names for CMYK subsets	PostScript transform function	DeviceN colorant names for CMYK subsets	PostScript transform function
Cyan Magenta	{0 0}	Cyan Yellow Black	{0 3 1 roll}
Cyan Yellow	{0 exch 0}	Magenta Yellow Black	{0 4 1 roll}
Cyan Black	{0 0 3 -1 roll}	Cyan Magenta Yellow Black	{ }
Magenta Yellow	{0 0 4 1 roll}		

A similar technique can be used to render only a subset of CMYK image channels. For example, the following code fragment creates a DeviceN color space which uses only the Black channel and replaces the other three channels with *None*. Since four color channels are present in the image data, the PostScript transform function replaces the Cyan, Magenta and Yellow channels with zero values (this differs from the PostScript functions in Table 4.4 which add missing 0 values instead of replacing existing color values with zero):

```
devicen = p.create_devicen(
    "names={None None None Black} " +
    "alternate=devicecmymk transform={{4 1 roll pop pop pop 0 0 0 4 -1 roll}}");
optlist = "width=4000 height=3000 bpc=8 colorize=" + devicen;
image = p.load_image("raw", filename, optlist);
```

Cookbook A full code sample can be found in the *Cookbook* topic `color/`
`colorize_image_with_DeviceN_color`.

DeviceN color space based on spot colors with Lab alternate space. Using two or more spot colors as DeviceN colorants provides the basis for shadings between spot colors. Implementing a suitable PostScript transform function for device-dependent alternate color spaces is rather challenging due to the nonlinear characteristics of RGB and CMYK. Blending colors in the Lab color space is much easier because it is perceptually uniform. This means we can combine colors simply by calculating the weighted average of their Lab values. The following PostScript code implements this method for two Pantone colors whose Lab alternate values must be provided at the start of the code:

```
% PostScript transform function for DeviceN with N=2 and Lab alternate;
% blend Lab colors values by using the tint values as
% weights and calculating the weighted average of the L/a/b values

80 28 75                % Lab values of color 1=PANTONE 123 U
31.7647 0 -17          % Lab values of color 2=PANTONE 289 U

% blend L values
7 index 6 index mul    % t1*L1
7 index 4 index mul    % t2*L2
add 9 1 roll           % bottom: L = t1*L1 + t2*L2

% blend a values
7 index 5 index mul    % t1*a1
7 index 3 index mul    % t2*a2
add 9 1 roll           % bottom: a = t1*a1 + t2*a2
```

```
% blend b values
7 index 4 index mul      % t1*b1
7 index 2 index mul      % t2*b2
add 9 1 roll              % bottom: b = t1*b1 + t2*b2
```

```
% pop 2 tint and 2x3 input color values
pop pop pop pop pop pop pop
% result: Lab values of blended color
```

The PostScript function *transformFunc2* above can be used to create a DeviceN color space based on the two Pantone spot colors as follows:

```
devicen = p.create_devicen(
  "names={{PANTONE 123 U} {PANTONE 289 U}} alternate=lab " +
  "transform={{" + transformFunc2 + "}}");
```

PDFlib automatically constructs such a DeviceN color space when a shading is constructed based on different spot colors with Lab alternate values.

Cookbook An extensive code sample and Lab-based PostScript transform functions for higher values of *N* can be found in the *Cookbook* topic `color/devicen_color`.

4.6 Shadings and Shading Patterns

Cookbook Full code samples can be found in the `starter_color` sample and the *Cookbook* topic `color/color_gradient`.

Smooth shadings, also called color blends or gradients, provide a continuous transition between two or more colors in the same color space, e.g. two RGB colors or two tints of a spot color. Simple shadings define a transition between two colors. The first color is taken from the *startcolor* option of *PDF_shading()* or the current fill color, The second color is provided in the *endcolor* option or the *c1*, *c2*, *c3*, and *c4* parameters. Shadings with an arbitrary number of intermediate colors can be created by using the *stopcolors* option instead of *startcolor/endcolor*.

The shading color is controlled by a variable *t* which varies linearly from 0 at the start color to 1 at the end color. The linear transition can be changed to an exponential transition with the *N* option as exponent. Since the RGB and CMYK color spaces are not perceptually linear, the resulting blends may not appear as smooth as desired. Specifying the blend with start and end colors in the Lab color space may result in a smoother color transition. PDFlib supports two different kinds of geometry for shadings:

- ▶ The color in axial shadings (*type=axial*) varies along a line between a starting point and an ending point (see figure 4.2). The shading color varies linearly from the starting point (*x0*, *y0*) to the ending point (*x1*, *y1*). The shading can optionally be extended

with the boundary color beyond the starting and/or ending point, subject to the *extendo* and *extend1* options.

- ▶ The color in radial shadings (*type=radial*) varies between two circles (see figure 4.3). Radial shadings may be used to create a 3D-like visualization of a sphere. The shading color varies from a starting circle with center (x_0, y_0) and radius r_0 to an ending circle with center (x_1, y_1) and radius r_1 . One of the circles may collapse to a point.

PDF_shading() returns a handle to a shading object which can be used in two ways:

- ▶ Fill an area directly with *PDF_shfill()*. This method is preferable if the geometry of the object to be filled is the same as the geometry of the shading. Contrary to its name this method not only fills the interior of the object, but also affects the exterior. This behavior can be modified with *PDF_clip()*.
- ▶ Define a shading pattern (not to be confused with tiling patterns) to be used for filling more complex objects. This involves calling *PDF_shading_pattern()* to create a pattern based on the shading and using this pattern to fill or stroke arbitrary objects.

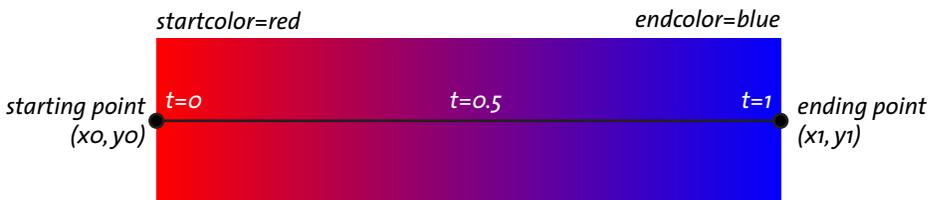
Shadings between two process colors. The following code creates an axial shading from red to blue in the RGB color space and uses it to fill a circle:

```
sh = p.shading("axial", 100, 100, 500, 500, 0, 0, 0, 0, "startcolor=red endcolor=blue");
shp = p.shading_pattern(sh, "");
p.set_graphics_option("fillcolor={pattern " + shp + "}");

p.circle(300, 300, 200);
p.fill();
```

Shadings between spot colors. Since shadings can only be used to create transitions between colors in the same color space, you cannot use arbitrary spot colors or a spot color and a process color as start and end color. However, for spot colors with Lab alternate values shadings can be achieved with a suitable DeviceN color space similar to the one shown in the section »DeviceN color space based on spot colors with Lab alternate space«, page 84. PDFlib automatically creates such a DeviceN color space for spot colors

Fig. 4.2
Main parameters of an axial shading between two colors



as stop colors, provided certain conditions are met (see PDFlib API Reference). The following code fragment creates a shading between two spot colors:

```
sh = p.shading("axial", 100, 100, 500, 500, 0, 0, 0, 0,
    "stopcolors={ 0% {spotname {PANTONE 123 U} 1} 100% {spotname {PANTONE 289 U} 1}}")
shp = p.shading_pattern(sh, "");
p.set_graphics_option("fillcolor={pattern " + shp + "}");

p.circle(300, 300, 200);
p.fill();
```

4.7 Tiling Patterns

Cookbook Full code samples can be found in the *starter_color* sample and the *Cookbook* topics *graphics/fill_pattern* and *images/tiling_pattern*.

A tiling pattern is defined by an arbitrary number of painting operations which are grouped into a single entity. This group of objects can be used to fill or stroke arbitrary other objects by replicating (or tiling) the group over the entire area to be filled or the path to be stroked. Working with patterns involves the following steps:

- ▶ First the pattern must be defined with drawing operators between *PDF_begin_pattern_ext()* and *PDF_end_pattern()*. Most graphics operators can be used to define a pattern.
- ▶ The pattern handle returned by *PDF_begin_pattern_ext()* can be used to set the pattern as the current color with the options *fillcolor/strokecolor* in *PDF_set_graphics_option()* or *PDF_setcolor()*.

Depending on the *painttype* option of *PDF_begin_pattern_ext()* the pattern definition may or may not include its own color specifications. If *painttype=colored*, the pattern definition must contain its own color specification and will always look the same; if *painttype=uncolored*, the pattern definition must not include any color specification. In-

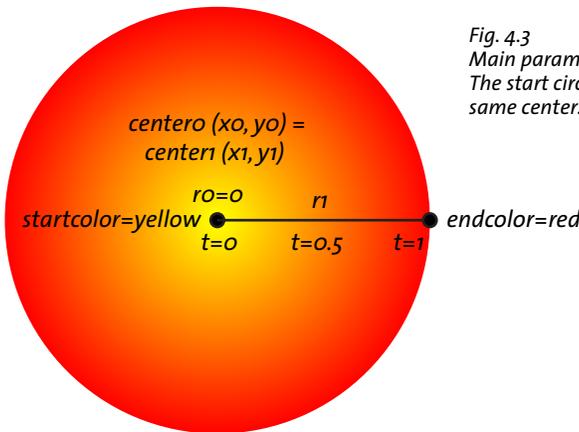


Fig. 4.3
Main parameters of a radial shading between two colors. The start circle collapses to a point; both circles have the same center.

stead, the current fill or stroke color will be applied when the pattern is used for filling or stroking.

4.8 Transparency Blend Modes

Cookbook A code sample can be found in the *Cookbook* topic `color/blendmode`.

The blend mode controls how the colors of an object blend with the background in the transparent imaging model (as opposed to the opaque imaging model where each object completely obscures all underlying objects). Since the background color may have been created by blending multiple other objects, it is referred to as *backdrop*. The object being painted is referred to as *source*. Source and backdrop colors are blended against each other with a blend function. The blend modes listed in Table 4.7 create a variety of artistic effects and can also be leveraged for specialized applications (see Section 4.9.1, »Changing the Color with Blend Modes«, page 91). The available blend modes can be categorized as follows:

- ▶ The default blend mode *Normal* makes the source completely opaque, i.e. the source object completely replaces the background. Setting a blend mode other than *Normal* implicitly makes the source object transparent, regardless of any opacity settings or image transparency.
- ▶ The blend modes *Darken*, *Multiply* and *ColorBurn* create darkening effects.
- ▶ The blend modes *Lighten*, *Screen* and *ColorDodge* create lightening effects.
- ▶ The blend modes *Overlay*, *SoftLight* and *HardLight* increase the contrast by modifying light and dark areas.
- ▶ The differencing blend modes *Difference* and *Exclusion* subtract color values. They are the only modes where painting a white object on a white backdrop does not result in white.
- ▶ The blend modes *Hue*, *Saturation*, *Color* and *Luminosity* are defined by their effect on the dimensions of the HSL color representation. In the HSL model hue describes the perceived color (red and green are different hues), saturation describes how colorful a color appears (how much gray is mixed into the color), and luminosity describes how dark or light a color appears (white has maximum luminosity).

Table 4.7 Blend modes and their effect when blending a source object against background objects (backdrop)

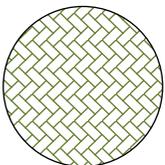
Blend mode	Description according to ISO 32000-2
------------	--------------------------------------

Blend modes without any blending effect

None	PDFlib-specific blend mode for creating a gstate without any blend mode specification. This ensures that a blend mode set in an enclosing gstate comes into effect.
-------------	---

Fig. 4.4

A tiling pattern used for filling a circle and some text. The text outlines are additionally stroked for emphasis.



Text filled with tiling pattern

Table 4.7 Blend modes and their effect when blending a source object against background objects (backdrop)

Blend mode	Description according to ISO 32000-2
Normal	Selects the source color, ignoring the backdrop.
Darkening blend modes	
Darken¹	Selects the darker of the backdrop and source colors. The backdrop is replaced with the source where the source is darker; otherwise, it is left unchanged.
Multiply¹	Multiplies the backdrop and source color values. The result color is always at least as dark as either of the two constituent colors. When working with additive colors, multiplying any color with black produces black while multiplying with white leaves the original color unchanged. For subtractive colors, the maximum tint value used for all colorants of the color space acts as black does for additive spaces. Painting successive overlapping objects with a color other than black or white produces progressively darker colors.
ColorBurn	Darkens the backdrop color to reflect the source color. Painting with white produces no change.
Lightening blend modes	
Lighten¹	Selects the lighter of the backdrop and source colors. The backdrop is replaced with the source where the source is lighter; otherwise, it is left unchanged.
Screen¹	Multiplies the complements of the backdrop and source color values, then complements the result. The result color is always at least as light as either of the two constituent colors. When working with additive colors, screening any color with white produces white while screening with black leaves the original color unchanged. For subtractive colors, the maximum tint value of all colorants of the color space acts as black does for additive spaces. The effect is similar to projecting multiple photographic slides simultaneously onto a single screen.
ColorDodge	Brightens the backdrop color to reflect the source color. Painting with black produces no change.
Contrasting blend modes	
HardLight	Multiplies or screens the colors if the source color value is smaller or larger than 0.5, respectively. The effect is similar to shining a harsh spotlight on the backdrop.
SoftLight	Darkens or lightens the colors if the source color value is smaller or larger than 0.5, respectively. The effect is similar to shining a diffused spotlight on the backdrop.
Overlay	Multiplies or screens the colors, depending on the backdrop color value. Source colors overlay the backdrop while preserving its highlights and shadows. The backdrop color is not replaced but is mixed with the source color to reflect the lightness or darkness of the backdrop.
Differencing blend modes	
Difference^{1,2}	Subtracts the darker of the two constituent colors from the lighter color: Painting with white inverts the backdrop color; painting with black produces no change. For subtractive colors, the maximum tint value for all colorants of the color space acts as black does for additive spaces.
Exclusion^{1,2}	Produces an effect similar to that of the Difference mode but lower in contrast. Painting with white inverts the backdrop color; painting with black produces no change. For subtractive colors, the maximum tint value for all colorants of the color space acts as black does for additive spaces.
HSL blend modes	
Hue²	Creates a color with the hue of the source color and the saturation and luminosity of the backdrop color.
Saturation²	Creates a color with the saturation of the source color and the hue and luminosity of the backdrop color. Painting with this mode in an area of the backdrop that is a pure gray (no saturation) produces no change.

Table 4.7 Blend modes and their effect when blending a source object against background objects (backdrop)

Blend mode	Description according to ISO 32000-2
Color²	Creates a color with the hue and saturation of the source color and the luminosity of the backdrop color. This preserves the gray levels of the backdrop and is useful for coloring monochrome images or tinting color images.
Luminosity²	Creates a color with the luminosity of the source color and the hue and saturation of the backdrop color. This produces the same effect as the Color mode, but with source and backdrop exchanged.

1. This mode is symmetric, i.e. swapping the source and backdrop colors doesn't change the result of the blending operation.
2. This mode doesn't have any effect on spot colors; blend mode Normal is used instead.

Blending color space. The PDF viewer performs all transparency calculations in a blending color space. This color space plays an important role since the blending results depend on the selected color space. It is determined as follows:

- ▶ The color space specified with the *colorspace* suboption of the *transparencygroup* page option if present.
- ▶ Otherwise the PDF/X or PDF/A output intent ICC profile is used if present. In PDF/X an output intent is always present. If a PDF/A document doesn't contain any output intent it must specify the *colorspace* suboption of the *transparencygroup* page option.
- ▶ Otherwise the native color space of the output device (or the simulated output device for soft proofing/output preview) is used. It is recommended to avoid this situation by specifying the blending color space explicitly since otherwise the rendering of transparent objects depends on the output device.

The following code fragment specifies DeviceRGB as blending color space and blend mode *Multiply*:

```
p.begin_page_ext(0, 0, "width=a4.width height=a4.height " +
    "transparencygroup={colorspace=DeviceRGB}");

gstate = p.create_gstate("blendmode=Multiply");
p.set_gstate(gstate);
```

Multiple blend modes. In PDF only one blend mode at a time can be active in the graphics state. If you want to apply the effects of multiple blend modes you can paint another object on top of the already blended backdrop and object. For example, in order to invert the result of some other blend mode, first blend the backdrop and object with the first blend mode; then change the blend mode to *Difference* and paint a white rectangle on top of the result of the first blend operation.

4.9 Changing the Color of Objects

PDFlib supports several methods for changing the color of objects such as imported PDF pages, raster images or SVG graphics, or arbitrary text and vector elements placed on the page:

- ▶ Raster images can be colorized, see Section 8.1.5, »Colorize Images with Spot or DeviceN Color«, page 168.
- ▶ The color of arbitrary objects can be changed with blend modes.
- ▶ GStates with luminosity soft masks can be used to apply color to arbitrary objects.

4.9.1 Changing the Color with Blend Modes

Cookbook A code sample for all effects discussed in this section can be found in the *Cookbook* topic `color/blendmode_effects`.

In this section we present some useful applications of blend modes. A description of all available blend modes can be found in Section 4.8, »Transparency Blend Modes«, page 88. In all examples the colors of an object, e.g. an imported image, PDF page, or SVG graphics, are modified in some way.

All examples can be implemented with the basic code fragment shown below. It specifies a transparency group color space for the page to avoid device-specific rendering and sets a specific blend mode. In order to avoid subsequent objects to be affected by the specified blend mode the sequence should be bracketed with a *save/restore* pair:

```
p.begin_page_ext(width, height, "transparencygroup={colorspace=DeviceCMYK}");

// place raster image, imported PDF page or SVG graphics
p.fit_image(image, 200, 150, optlist);

// Colorize, decolorize or invert existing content by blending with the desired color
p.save();
    // Create gstate with the desired blend mode
    gs_blendmode = p.create_gstate("blendmode=Color");
    p.set_graphics_option("fillcolor=red gstate=" + gs_blendmode);
    p.rect(0, 0, width, height);
    p.fill();
p.restore();
```

Colorize an object with arbitrary color. Blend mode *Color* can be used to colorize an object with some new color, i.e. the perceived gray levels of the object control the tint value of an additional color while the object's original colors are ignored. In order to colorize an object apply blend mode *Color* and place a colored rectangle (or other shape) on top of the object. In order to colorize transparent objects you it is recommend to place a white area below the placed object.

If you want the backdrop color to show through the white areas of the source object instead (while dark areas remain dark), use blend mode *Multiply*.

Decolorize or invert an object. In some situations you may want to decolorize (or desaturate, or bleach) an object, i.e. create grayscale output for an object which originally contained color, where the gray levels correspond to the perceived brightness (luminosity) of the object's original colors. This can be achieved with blend mode *Color* and a white area on top of the object.

The colors of raster images can be inverted with the *invert* image option. The colors of arbitrary objects can be inverted with blend mode *Difference* and white on top of the object.

4.9.2 Changing the Color with Soft Masks

Cookbook A code sample can be found in the *Cookbook* topic `color/softmask_effects`.

Image alpha channels can be generalized to arbitrary objects and color spaces with the help of a luminosity soft mask in a graphics state (gstate). The soft mask is constructed on the basis of a template with arbitrary contents such as imported PDF pages or SVG

graphics. The luminosity, i.e. perceived gray level or lightness of the template contents determines the visibility of subsequently drawn objects: light areas of the template are transparent (i.e. drawn objects are visible), while dark areas prevent objects from being visible. This technique entails the following steps:

- ▶ Create a template with the *transparencypgroup* option and suboption *colorspace*. Apply arbitrary drawing operations in the template.
- ▶ Create a gstate based on the template with the *softmask* option and the suboption *type=luminosity*. When applying the soft mask the light (white) areas of the soft mask template are colored with the current color while dark (black) areas don't get any color. The *invert* suboption of the *softmask* option can be used to reverse the effect of dark and light colors, i.e. dark and light areas are preserved. By default the soft mask background is initialized to opaque, i.e. the equivalent of all-black. This can be changed with the suboption *backdropcolor=transparent*, i.e. the equivalent of all-white. Note that the *invert* suboption also reverses the effect of *backdropcolor*.
- ▶ Set the gstate, usually bracketed with a *save/restore* sequence to limit the effect of the soft mask. The geometry of the soft mask template in the gstate is subject to the current coordinate system. For example, you can control the position of the soft mask by applying a translation before setting the gstate. Now draw the contents which will be masked on the page.

The following code fragment implements this sequence and applies it to colorize an imported SVG graphics file:

```
// Create the template which will be used to define the soft mask
tpl = p.begin_template_ext(595, 842, "transparencypgroup={colorspace=devicecmymk}");

    // Place arbitrary contents, e.g. SVG graphics
    p.fit_graphics(graphics, 0, 0, "boxsize={595 842} position={center} fitmethod=meet");

p.end_template_ext(0, 0);

// Create a gstate with inverted soft mask based on the template;
// The "invert" option makes black transparent, i.e. preserves dark and light areas
gstate_softmask = p.create_gstate(
    "softmask={type=luminosity invert template=" + tpl + " backdropcolor=transparent }");

// Activate the gstate with soft mask and apply color
// Result: SVG graphics is colorized with red
p.save();
    p.set_graphics_option("fillcolor=red gstate=" + gstate_softmask);
    p.rect(0, 0, 595, 842);
    p.fill();
p.restore();
```

Multiple soft masks. PDF supports only up to one soft mask per compositing operation. For example, if an image contains its own alpha channel (=soft mask), a soft mask supplied in a graphics state doesn't have any effect. As a workaround to this PDF restriction you must create the contents in a template and place this template on the page, effectively forcing the PDF viewer to apply a separate compositing operation where the soft mask can be used.

4.10 Overprint Control

Cookbook A code sample can be found in the *Cookbook* topic `color/overprint`.

By default, painting in a specific color channel replaces the corresponding areas of the other color channels. For example, color `o/o/1/o`, i.e. Yellow, on a CMYK device erases existing Cyan, Magenta and Black channels of objects which have been painted earlier at the same location of the page. This behavior can be changed with a process called overprinting which relates to mixing of color ink particles on a printing device. Overprinting can be used to create »enriched Black« by printing Black over another dark color.

The overprint flag in PDF is targeted at CMYK printing. It controls whether painting in a particular channel has an effect on existing contents in other channels. For example, if Cyan is printed on top of Yellow, the result will be green with overprinting (Cyan and Yellow are mixed). On the other hand, the same sequence without overprinting results in the last printed color Cyan since pure Cyan has no Yellow component, therefore `o% Yellow` replaces the existing Yellow contents.

Overprint settings for fill and stroke operations. The overprinting behavior can be controlled separately for stroke and other operations with the graphics appearance options `overprintstroke` (for stroking operations) and `overprintfill` (for all non-stroking operations including image placement). Both options default to `false`:

- ▶ If `overprintfill=false` or `overprintstroke=false`, filling or stroking in any color channel replaces the corresponding areas of unspecified colorants: the foreground color wins.
- ▶ If the options are `true` and the output device supports overprinting, previous markings in unspecified colorants are left unchanged. Overprinting is usually used for very dark colors.

Overprint mode for CMYK colors. The overprint behavior of CMYK objects can further be modified with the `overprintmode` graphics appearance option which controls the behavior of zero CMYK components if `overprintfill/overprintstroke=true`:

- ▶ `overprintmode=0` (default): each color component replaces previously placed marks (»foreground color wins«). In other words, all channels of the foreground object color knock out the underlying objects.
- ▶ `overprintmode=1` (sometimes called »Illustrator overprint mode«): tint value `0` leaves the corresponding component of previously painted color unchanged instead of setting it to `0` (»foreground tint value `0` is ignored«). In other words, only foreground channels with non-zero values knock out underlying objects. This is equivalent to painting in a DeviceN color space which includes only those CMYK components which have non-zero values.

This means that a small difference in CMYK tint values (e.g. `0%` vs. `1%`) can result in a big difference depending on the overprint mode: with overprint mode `0` the result is `0%` or `1%`, respectively. With overprint mode `1`, the result is the background color (whatever that was) or `1%`, respectively. Due to its definition overprint mode doesn't have any effect on color spaces other than CMYK.

Only colors in the DeviceCMYK color space are affected by overprinting. If the Cyan, Magenta, Yellow or Black components are used in a DeviceN color space the CMYK overprinting rules don't apply.

Device-dependent overprint rendering. Since PDF viewers must ignore overprint settings if the output device doesn't support overprinting, documents with overprinting objects are inherently non-portable and may require specific configuration of the PDF viewer or RIP to achieve the intended page representation. Whether or not Acrobat honors overprint settings for screen display can be configured via *Edit, Preferences, Page Display, Use Overprint Preview*. We recommend to set this preference to *Always* when working with overprint settings to ensure reliable screen display. In addition, the *Output Preview* panel in Acrobat also offers an item *Simulate Overprinting*.

If transparency or blend modes are used on the page, Acrobat's overprint preview works as expected only if the current blending color space is set to use the DeviceCMYK color space, e.g. with the following option in *PDF_begin_page_ext()*:

```
transparencygroup={colorspace=devicecmymk}
```

Third-party PDF viewers without support for overprint preview may display unexpected colors. In fact, most simple PDF viewers don't support overprint preview.

Device-independent overprinting effect. Because of the viewer dependencies described above the rendering of overprinting may be fragile. You can achieve overprinting effects in a device-independent manner by using blend mode *Darken* as follows:

```
gstate = p.create_gstate("blendmode=darken");  
p.set_gstate(gstate);
```

A particular situation where overprint simulation is useful is the recombination of color separations. Assuming you have the Cyan, Magenta, Yellow and Black separations of an image or page and need to recombine these into a composite PDF page, this can be achieved with the *Darken* blend mode.

Cookbook A code sample for recombining CMYK separations can be found in the *Cookbook* topic *color/recombine_color_channels*.

5 Unicode and Legacy Encodings

This chapter provides basic information about Unicode and other encoding schemes. Text handling in PDFlib heavily relies on the Unicode standard, but also supports various legacy encodings.

5.1 Important Unicode Concepts

Characters and glyphs. When dealing with text it is important to clearly distinguish the following concepts:

- ▶ *Characters* are the smallest units which convey information in a language. Common examples are the letters in the Latin alphabet, Chinese ideographs, and Japanese syllables. Characters have a meaning: they are semantic entities.
- ▶ *Glyphs* are different graphical variants which represent one or more particular characters. Glyphs have an appearance: they are representational entities.

There is no one-to-one relationship between characters and glyphs. For example, a ligature is a single glyph which represents two or more separate characters. On the other hand, a specific glyph may represent different characters depending on the context (some characters look identical, see figure 5.1).

BMP and PUA. The following terms will occur frequently in Unicode-based environments:

- ▶ The *Basic Multilingual Plane (BMP)* comprises the code points in the Unicode range U+0000...U+FFFF. The Unicode standard contains many more code points in the supplementary planes, i.e. in the range U+10000...U+10FFFF.
- ▶ The *Private Use Area (PUA)* consists of multiple Unicode ranges which are reserved for private use. PUA code points cannot be used for general interchange since the Unicode standard does not specify any characters in this range. The Basic Multilingual Plane includes the PUA range U+E000...U+F8FF. Plane fifteen (U+F0000... U+FFFFFD) and plane sixteen (U+100000...U+10FFFD) are completely reserved for private use.

Characters

U+0067 LATIN SMALL LETTER G

U+0066 LATIN SMALL LETTER F +
U+0069 LATIN SMALL LETTER I

U+2126 OHM SIGN or
U+03A9 GREEK CAPITAL LETTER OMEGA

U+2167 ROMAN NUMERAL EIGHT or
U+0056 V U+0049 I U+0049 I

Glyphs

Fig. 5.1
Relationship of glyphs
and characters

Unicode encoding forms (UTF formats). The Unicode standard assigns a number (code point) to each character. In order to use these numbers in computing, they must be represented in some way. In the Unicode standard this is called an encoding form (formerly: transformation format); this term should not be confused with font encodings. Unicode defines the following encoding forms:

- ▶ **UTF-8:** This is a variable-width format where code points are represented by 1-4 bytes. ASCII characters in the range U+0000...U+007F are represented by a single byte in the range 00...7F. Latin-1 characters in the range U+00A0...U+00FF are represented by two bytes, where the first byte is always 0xC2 or 0xC3 (these values represent Á and Ã in Latin-1).
- ▶ **UTF-16:** Code points in the Basic Multilingual Plane (BMP) are represented by a single 16-bit value. Code points in the supplementary planes, i.e. in the range U+10000...U+10FFFF, are represented by a pair of 16-bit values. Such pairs are called surrogate pairs. A surrogate pair consists of a high-surrogate value in the range D800...DBFF and a low-surrogate value in the range DC00...DFFF. High- and low-surrogate values can only appear as parts of surrogate pairs, but not in any other context.
- ▶ **UTF-32:** Each code point is represented by a single 32-bit value.

Unicode encoding schemes and the Byte Order Mark (BOM). Computer architectures differ in the ordering of bytes, i.e. whether the bytes constituting a larger value (16- or 32-bit) are stored with the most significant byte first (big-endian) or the least significant byte first (little-endian). A common example for big-endian architectures is PowerPC, while the x86/x64 architecture is little-endian. Since UTF-8 and UTF-16 are based on values which are larger than a single byte, the byte-ordering issue comes into play here. An encoding scheme (note the difference to encoding form above) specifies the encoding form plus the byte ordering. For example, UTF-16BE stands for UTF-16 with big-endian byte ordering. If the byte ordering is not known in advance it can be specified by means of the code point U+FEFF, which is called Byte Order Mark (BOM). Although a BOM is not required in UTF-8, it may be present as well, and can be used to identify a stream of bytes as UTF-8. Table 5.1 lists the representation of the BOM for various encoding forms.

Table 5.1 Byte order marks for various Unicode encoding forms

encoding form	byte order mark (hex)	graphical representation in WinAnsi ¹
UTF-8	EF BB BF	ï»¿
UTF-16 big-endian	FE FF	þÿ
UTF-16 little-endian	FF FE	ÿþ
UTF-32 big-endian	00 00 FE FF	■ ■ þÿ
UTF-32 little-endian	FF FE 00 00	ÿþ ■ ■

1. The square ■ denotes a null byte.

5.2 Unicode-capable Language Bindings

Some aspects of the PDFlib API vary depending on whether or not the used language binding is Unicode-capable. This concept is discussed in this section and the next section.

5.2.1 Language Bindings with native Unicode Strings

If a programming language or environment supports Unicode strings natively we call the binding Unicode-capable. The following PDFlib language bindings are Unicode-capable:

- ▶ C++
- ▶ .NET
- ▶ Java
- ▶ Objective-C
- ▶ Python
- ▶ RPG

String handling in these environments is straightforward: all strings are supplied to the PDFlib kernel as Unicode strings in native UTF-16 format. The language wrappers correctly deal with Unicode strings provided by the client and automatically set certain PDFlib options with the following consequences:

- ▶ All strings supplied by the client arrive in PDFlib in Unicode encoding and UTF-16 format.
- ▶ The distinction between different string types (content strings, hypertext strings and name strings) in the API descriptions is not relevant. The options *textformat*, *hypertextformat* and *hypertextencoding* are not required and are not allowed. The Textflow option *fixedtextformat* is forced to *true*.
- ▶ Using *unicode* encoding for the contents of a page is the easiest way to deal with encodings in Unicode-capable languages, but 8-bit encodings and single-byte text for symbol fonts can also be used if desired.

The overall effect is that clients can provide native Unicode strings to PDFlib API methods without any additional configuration.

5.2.2 Language Bindings with UTF-8 Support

Programming languages without native Unicode string data type can nevertheless deal with Unicode strings in UTF-8 format. All strings at the API, i.e. name strings, content strings, hypertext strings and option lists are expected in UTF-8 format with or without BOM.

If you must deal with strings in other encodings than Unicode, you must convert them to Unicode in UTF-8 before passing them to PDFlib. This can be achieved with *PDF_convert_to_unicode()* or language-specific methods.

Language bindings with default UTF-8 support. The following PDFlib language bindings are Unicode-capable by default via UTF-8 strings:

- ▶ Perl
- ▶ PHP
- ▶ Ruby

C binding with optional UTF-8 support. By default the C PDFlib language binding is not Unicode-capable (see below), but can be configured for use with UTF-8 strings instead:

```
PDF_set_option(p, "stringformat=utf8");
```

The format UTF-8 with BOM has the advantage for C users that PDFlib automatically recognizes such strings via the BOM. This makes it possible to load fonts with *encoding=unicode*, treat hypertext strings with *hypertextencoding=unicode* and name strings with *usehypertextencoding=true*, resulting in a full Unicode workflow.

The C interface still interprets strings passed as function parameters as UTF-16 if the *length* parameter is supplied with a value larger than 0.

5.3 Non-Unicode-capable Language Bindings

5.3.1 Language Bindings without native Unicode Strings

Language bindings with legacy string support. The following PDFlib language bindings use UTF-8 by default (see above), but can be configured for use with legacy 8-bit-encoded strings instead by setting *stringformat=legacy*:

- ▶ Perl
- ▶ PHP
- ▶ Ruby

In order to use legacy 8-bit-encoded strings with these language bindings you must set the string format as follows:

```
p.set_option("stringformat=legacy");
```

C binding with default support for legacy strings. The C binding works with 8-bit-encoded strings via the default setting *stringformat=legacy*.

The C interface still interprets strings passed as function parameters as UTF-16 if the *length* parameter is supplied with a value larger than 0.

5.3.2 Unicode Handling and String Types

Unicode strings can still be used in non-Unicode-capable languages, but string handling is more complicated and depends on the intended use of the string. The PDFlib API uses the string types *Content String*, *Hypertext String*, and *Name String* (these designations are historical misnomers). Parameters and options are marked as one of these types in the PDFlib API Reference. Handling of these string types is summarized in Table 5.2 and detailed below.

Content strings. Content strings are used to create page content according to the encoding selected for a particular font. All *method* parameters with the name *text* in the PDFlib API Reference for page content functions fall in this class. Since content strings are represented with glyphs from a particular font, the range of usable characters depends on the font/encoding combination.

Interpretation of content strings is controlled by the *textformat* option (detailed below) and the *selected* encoding. If *textformat=auto* (which is the default) *utf16* format is used for the *unicode* and *glyphid* encodings. For all other encodings the format is *bytes*.

In the C language binding the length of UTF-16 strings must be supplied in a separate *length* parameter.

Table 5.2 Summary of string handling for different string types

string type	sample parameters and options	relevant options/interpretation
content string	The text parameter of <code>PDF_fit_textline()</code> and <code>PDF_add_textflow()</code> .	textformat encoding
hypertext string	<ul style="list-style-type: none"> ▶ The fieldname option of <code>PDF_add_table_cell()</code> ▶ The name option of <code>PDF_define_layer()</code> ▶ The destname option of <code>PDF_create_action()</code> ▶ The text parameter of <code>PDF_create_bookmark()</code> 	hypertextformat hypertextencoding
name string	<ul style="list-style-type: none"> ▶ The filename parameter of <code>PDF_begin_document()</code> and <code>PDF_create_pvf()</code> ▶ The fontname parameter of <code>PDF_load_font()</code> ▶ The profilename parameter of <code>PDF_load_iccprofile()</code> 	usehypertextencoding, filenamehandling
strings in option lists		with BOM: UTF-8 without BOM: depends on string type

Hypertext strings. Hypertext strings are used for interactive features such as bookmarks and annotations, and are labeled *Hypertext string* in the PDFlib API Reference. Many parameters and options of the functions for interactive features fall in this class.

Interpretation of hypertext strings is controlled by the *hypertextformat* and *hypertextencoding* options (detailed below). If *hypertextformat=auto* (which is the default) *utf16* format is used if *hypertextencoding=unicode*, and *bytes* otherwise.

In the C language the length of UTF-16 strings must be supplied in a separate *length* parameter.

Name strings. Name strings are used for external file names, font names, Block names, etc., and are marked as *Name string* in the PDFlib API Reference. They slightly differ from Hypertext strings.

File names are a special case: the option *filenamehandling* specifies how PDFlib converts filenames supplied to the API to a string which can be used with the local file system.

By default, name strings are interpreted in *host* encoding. However, if a name starts with a UTF-8 BOM it is interpreted as UTF-8 (or as EBCDIC UTF-8 if it starts with an EBCDIC UTF-8 BOM). If *usehypertextencoding=true*, the encoding specified in *hypertextencoding* is applied to name strings as well. This can be used, for example, to specify font or file names in Shift-JIS. If *hypertextencoding=unicode* PDFlib expects a UTF-16 string which must be terminated by two null bytes.

In C the *length* parameter must be 0 for UTF-8 strings. If it is different from 0 the string is interpreted as UTF-16. In other non-Unicode-capable language bindings name strings must always be supplied in UTF-8 format.

Text format for content strings and hypertext strings. Unicode strings can be supplied in UTF-8 or UTF-16 format with any byte ordering. The choice of format can be controlled with the *textformat* option for text on page descriptions and the *hypertextformat* option for interactive elements. Table 5.3 lists the values which are supported for these options. The default for the *[hyper]textformat* option is *auto*. Use the *usehypertextencoding* option to enforce the same behavior for name strings. The default for the *hypertextencoding* option is *auto*.

Table 5.3 Values for the `textformat` and `hypertextformat` options

<code>[hyper]textformat</code>	explanation
<code>bytes</code>	One byte in the string corresponds to one character. This is useful for 8-bit encodings and symbolic fonts. A UTF-8 BOM at the start of the string will be evaluated and then removed.
<code>utf8</code>	Strings are expected in UTF-8 format. Invalid UTF-8 sequences trigger an exception if <code>glyphcheck=error</code> , and are deleted otherwise.
<code>ebcdicutf8</code>	Strings are expected in EBCDIC-coded UTF-8 format (only on IBM System i and IBM Z).
<code>utf16</code>	Strings are expected in UTF-16 format. A Unicode Byte Order Mark (BOM) at the start of the string is evaluated and then removed. If no BOM is present the string is expected in the machine's native byte ordering (e.g. little-endian on Intel x86/x64).
<code>utf16be</code>	Strings are expected in UTF-16 format in big-endian byte ordering. There is no special treatment for Byte Order Marks.
<code>utf16le</code>	Strings are expected in UTF-16 format in little-endian byte ordering. There is no special treatment for Byte Order Marks.
<code>auto</code>	<p>Content strings: equivalent to <code>bytes</code> for 8-bit encodings, and <code>utf16</code> for unicode and <code>glyphid</code> encoding.</p> <p>Hypertext strings: UTF-8 and UTF-16 strings with BOM will be detected (in C UTF-16 strings must be terminated with a double-null). If the string does not start with a BOM, it is interpreted as an 8-bit encoded string according to the <code>hypertextencoding</code> option.</p> <p>This setting provides proper text interpretation in most environments which do not use Unicode natively.</p>

Although the `textformat` setting is in effect for all encodings, it is most useful for *unicode* encoding. Table 5.4 details the interpretation of text strings for various combinations of encodings and `textformat`. If a code or Unicode value in a content string cannot be represented with a suitable glyph in the selected font, the option `glyphcheck` controls the behavior of PDFlib (see »Glyph replacement«, page 157).

Table 5.4 Relationship of encodings and text format

<code>[hypertext]encoding</code>	<code>textformat=bytes</code>	<code>textformat=utf8, utf16, utf16be or utf16le</code>
All string types:		
<code>auto</code>	see section »Automatic encoding«, page 102	
<code>unicode</code>	8-bit codes are Unicode values from U+0000 to U+00FF	any Unicode value, encoded according to the chosen text format ¹
Only content strings:		
8-bit encodings	8-bit codes	Convert Unicode values to 8-bit codes according to the chosen encoding ¹ .
<code>glyphid</code>	8-bit codes are glyph ids from 0 to 255	Unicode values are interpreted as glyph ids ² .

1. If the Unicode character is not available in the font, PDFlib throws an exception or replaces it subject to the `glyphcheck` option.
2. If the glyph id is not available in the font, PDFlib throws an exception or replaces it with glyph id 0 subject to the `glyphcheck` option.

Option lists. Strings within option lists require special attention since in non-Unicode-capable language bindings they cannot be expressed as Unicode strings in UTF-16 format, but only as byte strings. For this reason UTF-8 is used for Unicode options. By looking for a BOM at the beginning of an option, PDFlib decides how to interpret it. The BOM is used to determine the format of the string, and the string type (content string, hypertext string, or name string) is used to determine the encoding. More precisely, interpreting a string option works as follows:

- ▶ If the option starts with a UTF-8 BOM (*0xEF 0xBB 0xBF*) it is interpreted as UTF-8. On EBCDIC-based systems: if the option starts with an EBCDIC UTF-8 BOM (*0x57 0x8B 0xAB*) it is interpreted as EBCDIC UTF-8.
- ▶ If no BOM is found, string interpretation depends on the type of string:
 - ▶ Content strings are interpreted according to the applicable *encoding* option or the encoding of the corresponding font (whichever is present).
 - ▶ Hypertext strings are interpreted according to the *hypertextencoding* option.
- ▶ Name strings are interpreted according to the *hypertext* settings if *usehypertext-encoding=true*, and *auto* encoding otherwise.

The characters { and } require special handling within strings in option lists and must be preceded by a \ character if they are used within a string option.

5.3.3 Single-Byte (8-Bit) Encodings

8-bit encodings (map a byte value to a single character with a Unicode value in the BMP). They are limited to 255 different characters since code 0 (zero) is reserved for the *.notdef* character. PDFlib includes the following encodings internally:

```
winansi (identical to cp1252; superset of iso8859-1),
cp1252
iso8859-1
macroman (the original Macintosh character set),
ebcdic (EBCDIC code page 1047),
ebcdic_37 (EBCDIC code page 037),
pdfdoc (PDFDocEncoding)
```

The following encodings are available as separate files within the PDFlib packages:

```
iso8859-2, iso8859-3, iso8859-4, iso8859-5, iso8859-6, iso8859-7, iso8859-8, iso8859-9,
iso8859-10, iso8859-13, iso8859-14, iso8859-15, iso8859-16,
cp1250, cp1251, cp1253, cp1254, cp1255, cp1256, cp1257, cp1258
```

Unless they are available as system codepages, access to the resource files for these encoding definitions must be configured (see Section 3.1.4, »Resource Configuration and File Search«, page 51), e.g.

```
p.set_option("SearchPath={{/usr/pdflib/resource/codepage}}");
p.set_option("Encoding={iso8859-5=iso8859-5.cpg}");
```

Host encoding. The special encoding *host* does not have any fixed meaning, but is mapped to another 8-bit encoding depending on the current platform as follows:

- ▶ on IBM System Z with MVS or USS it is mapped to *ebcdic*;
- ▶ on IBM System i it is mapped to *ebcdic_37*;
- ▶ on Windows it is mapped to *winansi*;
- ▶ on all other systems it is mapped to *iso8859-1*;

Host encoding is primarily useful for writing platform-independent test programs and other simple applications. Host encoding is not recommended for production use, but should be replaced by whatever encoding is appropriate.

Automatic encoding. PDFlib supports a mechanism which can be used to specify the most natural encoding for certain environments. Supplying the keyword *auto* as encoding name selects a platform- and environment-specific 8-bit encoding for text fonts as follows:

- ▶ On Windows: the current system code page (see below for details)
- ▶ On Unix and macOS: *iso8859-1*
- ▶ On IBM System i: the current job's encoding (*IBMCCSID000000000000*)
- ▶ On IBM System Z: *ebcdic* (=code page 1047).

For symbol fonts the keyword *auto* is mapped to *builtin* encoding (see Section 6.3.2, »Selecting an Encoding for symbolic Fonts«, page 117). While automatic encoding is convenient in many circumstances, using this method makes your PDFlib client programs inherently non-portable.

Encoding *auto* is used as the default encoding for Name strings (see Section 5.3, »Non-Unicode-capable Language Bindings«, page 98) in non-Unicode-capable language bindings.

Tapping system code pages. PDFlib can fetch code page definitions from the system. Instead of supplying the name of a built-in or user-defined encoding for *PDF_load_font()*, simply use an encoding name which is known to the system. This feature is only available on selected platforms, and the syntax for the encoding string is platform-specific:

- ▶ On Windows the encoding name is *cp<number>*, where *<number>* is the number of any single-byte code page installed on the system (see Section 7.6.1, »Using TrueType and OpenType CJK Fonts«, page 150, for information on multi-byte Windows code pages):

```
font = p.load_font("Helvetica", "cp1250", "");
```

Single-byte code pages will be transformed into an internal 8-bit encoding, while multi-byte code pages will be mapped to Unicode at runtime. The text must be supplied in a format which is compatible with the chosen code page (e.g. SJIS for *cp932*, see »Code pages for CJK fonts«, page 104).

- ▶ On Linux all codeset identifiers supported the *iconv* facility can be used.
- ▶ On IBM System i any *Coded Character Set Identifier* (CCSID) can be used. The CCSID must be supplied as a string, and PDFlib will apply the prefix *IBMCCSID* to the supplied code page number. PDFlib will also add leading 0 characters if the code page number uses fewer than 5 characters. Supplying 0 (zero) as the code page number will result in the current job's encoding to be used:

```
font = p.load_font("Helvetica", "273", "");
```

- ▶ On IBM System Z with USS or MVS any *Coded Character Set Identifier* (CCSID) can be used. The CCSID must be supplied as a string, and PDFlib will pass the supplied code page name to the system literally without applying any change:

```
font = p.load_font("Helvetica", "IBM-273", "");
```

User-defined 8-bit encodings. In addition to predefined encodings PDFlib supports user-defined 8-bit encodings. These are the way to go if you want to deal with some character set which is not available in PDFlib. PDFlib supports encoding tables defined by glyph names as well as tables defined by Unicode values.

The following tasks must be done before a user-defined encoding can be used in a PDFlib program (alternatively the encoding can be constructed at runtime using `PDF_encoding_set_char()`):

- ▶ Generate a description of the encoding in a simple text format.
- ▶ Configure the encoding as PDFlib resource (see Section 3.1.4, »Resource Configuration and File Search«, page 51).
- ▶ Provide a font with glyphs for all characters used in the encoding.

The encoding file lists glyph names and codes line by line. The following excerpt shows the start of an encoding definition:

```
% Encoding definition for PDFlib, based on glyph names
% name      code   Unicode (optional)
space      32     0x0020
exclam     33     0x0021
...
```

If no Unicode value has been specified PDFlib searches for a suitable Unicode value in its internal tables. A Unicode value can be specified instead of a glyph name:

```
% Code page definition for PDFlib, based on Unicode values
% Unicode   code
0x0020     32
0x0021     33
...
```

▶

All Unicode values in an encoding or codepage file must be smaller than U+FFFF.

5.3.4 Chinese, Japanese, and Korean CMaps (deprecated)

Note All features for legacy encodings discussed in this section are deprecated.

Note The following features are not available for fonts loaded with a CMap: Textflow, glyph replacement, fallback fonts, shaping and OpenType layout features.

PDFlib supports CJK legacy encoding via the CMaps listed in Table 5.5. Legacy encodings can only be used in non-Unicode-aware language bindings, i.e. C, or Perl, PHP and Ruby with `stringformat=legacy`.

Table 5.5 CMaps for Japanese, Chinese, and Korean text

locale and character collection	CMap name
Simplified Chinese (Adobe-GB1-5)	UniGB-UTF16-H/V ¹ , GB-EUC-H/V, GBpc-EUC-H/V, GBK-EUC-H/V, GBKp-EUC-H/V, GBK2K-H/V
Traditional Chinese (Adobe-CNS1-7)	UniCNS-UTF16-H/V ¹ , B5pc-H/V, HKscs-B5-H/V, ETen-B5-H/V, ETenms-B5-H/V, CNS-EUC-H/V
Japanese (Adobe-Japan1-7)	UniJIS-UTF16-H/V ¹ , 83pv-RKSJ-H, 9oms-RKSJ-H/V, 9omsp-RKSJ-H/V, 9opv-RKSJ-H, Add-RKSJ-H/V, EUC-H/V, Ext-RKSJ-H/V, H/V

Table 5.5 CMaps for Japanese, Chinese, and Korean text

locale and character collection	CMap name
Korean (Adobe-Korea1-2)	(Deprecated in PDF 2.0) UniKS-UTF16-H/V ¹ , KSC-EUC-H/V, KSCms-UHC-H/V, KSCms-UHC-HW-H/V, KSCpc-EUC-H/V
Korean (Adobe-KR-9)	(PDF 2.0; unsupported) UniAKR-UTF16-H

1. Requires PDF 1.5 or above.

CMap resource configuration. In order to create Chinese, Japanese, or Korean (CJK) text output with a CMap PDFlib requires the corresponding CMap file for processing the incoming text and mapping CJK encodings to Unicode. CMap files are available in a separate package. They should be installed as follows:

- ▶ On Windows the CMap files are found automatically if you place them in the *resource/cmap* directory within the PDFlib installation directory.
- ▶ On other systems you can place the CMap files in any convenient directory, and must manually configure the CMap files by setting the *SearchPath* at runtime:

```
p.set_option("SearchPath={{/path/to/resource/cmap}}");
```

As an alternative method for configuring access to the CJK CMap files you can set the *PDFLIBRESOURCEFILE* environment variable to point to a UPR configuration file which contains a suitable *SearchPath* definition.

Keeping native CJK legacy codes. If the *keepnative* font option is set to *true* for unembedded fonts, native legacy codes (e.g. Shift-JIS) according to the selected CMap are written to the PDF output; otherwise the text is converted to CID values. The advantage is that such fonts can be used for form fields without embedding.

If *keepnative=false* legacy code sequences are converted to CID values when PDF output is created. The advantage is that OpenType features and the Textflow formatter can be used.

Code pages for CJK fonts. PDFlib supports the code pages listed in Table 5.6. PDFlib on Windows additionally supports any CJK code page installed on the system.

Table 5.6 CJK code pages (must be used with *textformat=auto* or *textformat=bytes*)

locale	code page	format	character set
Simplified Chinese	cp936	GBK	GBK
Traditional Chinese	cp950	Big Five	Big Five with Microsoft extensions
Japanese	cp932	Shift-JIS	JIS X 0208:1997 with Microsoft extensions
Korean	cp949	UHC	KS X 1001:1992, remaining 8822 hangul as extension
	cp1361	Johab	Johab

5.4 Addressing Characters

Some environments require the programmer to write source code in 8-bit encodings (such as *winansi* or *ebcdic*). This makes it cumbersome to include isolated Unicode characters in 8-bit encoded text. In order to aid developers in this situation, PDFlib supports several auxiliary methods for expressing text.

5.4.1 Escape Sequences

PDFlib supports a method for incorporating arbitrary values within text strings via *escape sequences* (this is actually a misnomer; *backslash substitution* might be a better term). For example, the `\t` sequence in the default text of a text Block can be used to include tab characters which may not be possible by direct keyboard input. Similarly, escape sequences are useful for expressing codes for symbolic fonts, or in literal strings for language bindings where escape sequences are not available.

An escape sequence is an instruction to replace a sequence with a single byte value. The sequence starts with the code for the backslash character `'\'` in the current encoding of the string. Since the backslash has special meaning in many programming languages it may have to be duplicated if used literally in source code. The byte values resulting from substituting escape sequences are listed in Table 5.7; some differ between ASCII and EBCDIC platforms. Only byte values in the range 0-255 can be expressed with escape sequences.

Unlike some programming languages, escape sequences in PDFlib always have fixed length depending on their type. Therefore no terminating character is required for the sequence.

Table 5.7 Escape sequences for byte values

<i>sequence</i>	<i>length</i>	<i>macOS, Windows, Unix</i>	<i>EBCDIC platforms</i>	<i>common interpretation</i>
<code>\f</code>	2	<code>oC</code>	<code>oC</code>	<i>form feed</i>
<code>\n</code>	2	<code>oA</code>	<code>15/25</code>	<i>line feed</i>
<code>\r</code>	2	<code>oD</code>	<code>oD</code>	<i>carriage return</i>
<code>\t</code>	2	<code>o9</code>	<code>o5</code>	<i>horizontal tabulation</i>
<code>\v</code>	2	<code>oB</code>	<code>oB</code>	<i>line tabulation</i>
<code>\\</code>	2	<code>5C</code>	<code>Eo</code>	<i>backslash</i>
<code>\xNN</code>	4	<i>two hexadecimal digits specifying a byte value, e.g. <code>\xFF</code></i>		
<code>\NNN</code>	4	<i>three octal digits specifying a byte value, e.g. <code>\377</code></i>		

Escape sequences are not substituted by default. You must explicitly set the *escape-sequence* option to *true* to use escape sequences in strings. This can be done selectively in the functions where it is required, e.g. `PDF_fit_textline()`. Alternatively you can activate escape sequence expansion for all content strings (i.e. all text output operations) as follows:

```
p.set_text_option("escapesequence=true");
```

Escape sequence substitution can also be enabled globally with `PDF_set_option()`. This global option affects all subsequently used name strings, hypertext strings and content strings. Since the values of environment variables are treated as name strings they are also subject to escape sequence expansion. This may result in undesired behavior e.g. with Windows path or file names containing backslash characters. For this reason it is not recommended to activate escape sequence substitution globally with `PDF_set_option()`, but only selectively with the method where it is required (e.g. `PDF_fit_textline()`) or only for content strings with `PDF_set_text_option()`.

As an alternative to setting the `escapesequence` option you can substitute escape sequences in strings using `PDF_convert_to_unicode()`, using the same input and output encoding, e.g.

```
String s_plain = p.convert_to_unicode("utf16",
    s.character.getBytes("UTF-16"),
    "outputformat=utf16 escapesequence=true");
```

Escape sequences are evaluated after BOM detection, but before converting to the target format. If `textformat= utf16, utf16le` or `utf16be` escape sequences must be expressed as two byte values according to the selected format. Each character in the escape sequence will be represented by two bytes, where one byte has the value zero. If `textformat=utf8` the resulting code will not be converted to UTF-8.

If an escape sequence cannot be resolved (e.g. `\x` followed by invalid hex digits) an exception is thrown. For content strings the behavior is controlled by the `glyphcheck` and `errorpolicy` settings.

5.4.2 Character References

Cookbook A full code sample can be found in the *Cookbook* topic `fonts/character_references`.

A character reference is an instruction to replace a sequence with a Unicode value. The sequence starts with the code of the ampersand character '&' in the current encoding and ends with the semicolon character ';'. There are several methods available for expressing the target Unicode values:

HTML character references. PDFlib supports all character entity references defined in HTML 4.0. Numeric character references can be supplied in decimal or hexadecimal notation. The full list of HTML character references can be found at the following location:

www.w3.org/TR/REC-html40/charset.html#h-5.3

Examples:

<code>&shy;</code>	U+00AD soft hyphen
<code>&euro;</code>	U+20AC Euro glyph (entity name)
<code>&lt;</code>	U+003C less than sign
<code>&gt;</code>	U+003E greater than sign
<code>&amp;</code>	U+0026 ampersand sign
<code>&Alpha;</code>	U+0391 Greek Alpha

Numerical character references. Numerical character references for Unicode characters are also defined in HTML 4.0. They require the hash character '#' and a decimal or hexadecimal number, where hexadecimal numbers are introduced with a lower- or uppercase 'X' character. Examples:

­	U+00AD soft hyphen (decimal)
­	U+00AD soft hyphen
å	U+0229 letter a with small circle above (decimal)
å	U+00E5 letter a with small circle above (hexadecimal)
å	U+00E5 letter a with small circle above (hexadecimal)
€	U+20AC Euro glyph (hexadecimal)
😂	U+1F602 Face with Tears of Joy

PDFlib-specific entity names. PDFlib supports custom character entity references for the following groups of Unicode control characters:

- ▶ Control characters for overriding the default shaping behavior listed in Table 7.6.
- ▶ Control characters for overriding the default bidi formatting listed in Table 7.7.
- ▶ Control characters for Textflow line breaking and formatting listed in Table 9.1.

Examples:

&linefeed;	U+000A linefeed control character
&hortab;	U+0009 horizontal tab
&ZWNJ;	U+200C ZERO WIDTH NON-JOINER

Glyph name references. A font may contain glyphs which are not directly accessible because the corresponding Unicode values are not known in advance (since PDFlib assigns PUA values at runtime). As an alternative for addressing such glyphs, character references with glyph names can be used. These references are replaced with the corresponding Unicode values. Glyph names are drawn from the following sources:

- ▶ Common glyph names are searched in an internal list.
- ▶ Font-specific glyph names are searched in the current font. Character references of this class work only with content strings since they require a font.

In order to identify glyph name references the actual name requires a period character '.' after the ampersand character '&'. Examples:

&.three;	U+0033 common glyph name for the digit 3
&.mapleleaf;	(PUA unicode value) custom glyph name from Carta font
&.T.swash;	(PUA unicode value) second period character is part of the glyph name

Character references with glyph names are useful in the following scenarios:

- ▶ Character references with font-specific glyph names are useful in contents strings to select alternate character forms (e.g. swash characters) and glyphs without any specific Unicode semantics (symbols, icons, and ornaments). Note that tabular figures and many other features are more easily implemented with OpenType features (see Section 7.4, »OpenType Layout Features«, page 139) if supported by the font.
- ▶ Names from the Adobe Glyph List (including the *uniXXXX* and *uXXXX/uXXXXX/uXXXXXX* forms) plus various common »misnamed« glyph names are always accepted for content strings and hypertext strings.

Byte value references. Numerical values can also be supplied in character references which may be useful for addressing the glyphs in a symbol font. This variant requires an additional hash character '#' and a decimal or hexadecimal number, where hexadecimal numbers are introduced with a lower- or uppercase 'X' character. Example (assuming the Wingdings font):

&.#x9F;	bullet symbol in Wingdings font
&.#159;	bullet symbol in Wingdings font

Using character references. Character references are not substituted by default; you must set the *charref* option to *true* to use character references in content strings:

```
p.fit_textline("Price: 500&euro;", x, y, "charref=true");
```

Supplying *charref* as text option enables character references for all content strings:

```
p.set_text_option("charref=true font=" + font + " fontsize=24");  
p.fit_textline("Price: 500&euro;", x, y, "");
```

The *charref* option can also be set globally with *PDF_set_option()*. However, this is not recommended because it affects all name strings, hypertext strings and content strings, which may have undesired results.

As an alternative to setting the *charref* option you can substitute character references in strings using *PDF_convert_to_unicode()*, using the same input and output encoding, e.g.

```
String s_plain = p.convert_to_unicode("utf16",  
    s.character.getBytes("UTF-16"),  
    "outputformat=utf16 charref=true");
```

Additional notes on using character references:

- ▶ Character references can be used in all content strings, hypertext strings, and name strings. As an exception, font-specific glyph name references work only with contents strings as noted above.
- ▶ Character references are not substituted in text with *builtin* encoding. However, you can use character references for symbolic fonts by using *unicode* encoding.
- ▶ Character references are not substituted in option lists, but they are recognized in options with the *Unichar* data type; in this case the '&' and ';' decoration must be omitted. This recognition is always enabled; it is not subject to the *charref* option.
- ▶ In non-Unicode-capable language bindings character references must be expressed as two-byte values if *textformat=utf16*, *utf16be*, or *utf16le*. If *encoding=unicode* and *textformat=bytes* the character references must be expressed in ASCII (even on EBCDIC-based platforms).
- ▶ An ampersand character '&' is considered as start of a character reference if a subsequent semicolon ';' can be found and the characters between '&' and ';' are valid for a character reference. Otherwise the ampersand character '&' is left unchanged. In order to express an ampersand character '&' without starting a character reference it is recommended to use the character reference *&*.
- ▶ If a potential character reference sequence is invalid (e.g. *&#* followed by invalid decimal digits or '&' followed by an unknown entity name) the behavior depends on the option *glyphcheck*:
 - glyphcheck=none*: the sequence is kept as is;
 - glyphcheck=replace*: the sequence is replaced with the replacement character.
 - glyphcheck=error*: an error occurs and text processing is stopped. If *errorpolicy=exception* an exception is thrown.

6 Fonts

6.1 Font Formats

6.1.1 TrueType Fonts

TrueType file formats. PDFlib supports vector-based TrueType fonts in the following file formats:

- ▶ Windows and macOS TrueType fonts (*.*ttf*), including Western, symbolic, and CJK fonts;
- ▶ TrueType collections (*.*ttc*) with multiple fonts in a single file. TTC files are typically used for grouping CJK fonts, but also to package multiple members of a Western font family in a single file.
- ▶ End-user defined character (EUDC) fonts (*.*tte*) created with Microsoft's *eudcedit.exe* tool;



TrueType font names. If you are working with font files you can assign an arbitrary API name to the font (see »Sources of Font Data«, page 122). This name is used for loading the font and may differ from the font file name and the font's internal name. In the generated PDF the name of a TrueType font may differ from the name used in PDFlib (or Windows). This is normal and results from the fact that PDF uses the PostScript name of a TrueType font, which differs from its TrueType name (e.g., *TimesNewRomanPSMT* vs. *Times New Roman*).

6.1.2 OpenType Fonts

The OpenType font format (also called Open Font Format per ISO 14496-22) combines PostScript and TrueType technology. It is implemented as an extension of the TrueType file format. OpenType fonts may contain optional tables which can be used to enhance text output, e.g. ligatures and swash characters (see Section 7.4, »OpenType Layout Features«, page 139), as well as tables for complex script shaping (see Section 7.5, »Complex Script Output«, page 145).



OpenType fonts may contain glyph descriptions based on TrueType or PostScript. The PostScript flavor is also called *Compact Font Format* (CFF) or *Type 2*, and is usually used with the *.*otf* suffix.

- ▶ The CID (Character ID) architecture is used for CJK fonts. Modern CID fonts are packaged as OpenType *.*otf* fonts with PostScript outlines.
- ▶ OpenType Collections have been introduced with the OpenType 1.7 specification. They are similar to TrueType Collections in that they package multiple related OpenType fonts into a single combined file. OpenType Collections use the *.*ttc* or *.*otc* suffix..
- ▶ Variable fonts use a mechanism called OpenType font variations and can be used to package multiple font faces within the same font family, e.g. light, regular and bold, in a single font resource. This mechanism cannot be used since font variations are not supported in PDF. OpenType font variations with TrueType outlines can be loaded with PDFlib, but only the default font instance without variations is available.

OpenType font variations with PostScript outlines in a CFF2 table are rejected since they are not supported in PDF.

6.1.3 OpenType Color Fonts in SVG, COLR and sbix Flavors

Cookbook A full code sample can be found in the *Cookbook* topic `text_output/emoji_sequences`.

OpenType color fonts in various flavors became quite popular, especially for emoji characters. PDFlib supports the OpenType color font flavors (according to ISO 14496-22:2019) listed below.

Since OpenType color fonts are not supported directly in PDF, PDFlib creates a Type 3 font for the color glyphs. PDFlib processes OpenType color fonts in the OTF, TTF, TTC, OTC, WOFF or WOFF2 file formats as well as color fonts installed as host fonts. PDF output with text in color fonts remains fully searchable and selectable since color glyphs are not converted to raster images.

OpenType SVG fonts. This font flavor uses SVG glyph descriptions (Adobe/Mozilla). SVG glyphs are processed according to the description of SVG import (see Section 8.2, »SVG Graphics«, page 170). SVG glyphs are handled in the sRGB color space and may contain transparency. A sample font is the open source font *EmojiOne Color*.



OpenType COLR fonts. This font flavor uses COLR and CPAL tables (Microsoft) and construct color glyphs by layering one or more mono glyphs, where each layer is colored with an sRGB color value and optionally assigned a transparency value. A sample font is the Windows system font *Segoe UI Emoji*.

Auxiliary glyphs used only as layers in a COLR font are used for creating color glyphs, but are not usable themselves except with `encoding=glyphid` where all glyphs can be addressed. As usual, such auxiliary glyphs are only included in the font subset if they are actually used in the document.

OpenType sbix fonts. This font flavor uses glyph bitmaps in the *sbix* table (Apple). Such fonts contain one or more sets of bitmap glyph with different numbers of pixels per design unit. PDFlib uses the largest bitmap by default, but this can be changed with the `strikeselect` option. Color fonts with *sbix* table are handled in the sRGB color space and usually contain glyph bitmaps with a transparency mask. A sample font is the macOS system font *Apple Color Emoji* (the iOS system font of the same name cannot be used because it uses a proprietary bitmap compression scheme).

Color font processing modes. Depending on the intended use of a color font and its characteristics the `colormode` font option can be used to select a processing mode for color fonts. Table 6.1 compares color font modes. The color mode of a font cannot be changed; if the font is loaded again the `colormode` option is ignored and the color mode of the first load operation is used.

Table 6.1 Processing modes for OpenType color fonts (font option `colormode`)

	<code>colormode=combined</code> (default)	<code>colormode=ignoremono</code>	<code>colormode=ignorecolor</code>
description	single font with mono and color glyphs	mono glyphs are ignored and a pure color font is created	color glyphs are ignored and a pure mono font is created
advantages	Consistent handling of mono and color glyphs; identical glyph IDs for color and mono glyphs as in the original OpenType font. Even layer glyphs in COLR fonts can be used with <code>encoding=glyphid</code> .	unneded OpenType font with mono glyphs is skipped	no undesired color glyphs or transparency
use cases	simple use of mono and color glyphs	color font as fallback for another font	b/w workflows and avoidance of transparency
restrictions	–	mono glyphs, OpenType features, ligatures and shaping n/a	color glyphs n/a

Glyph types in an OpenType color font. Color fonts may contain the following types of glyphs:

- ▶ Color glyphs, i.e. glyphs with intrinsic color and optionally transparency information. Since the color information is part of the glyph description the color is fixed for a particular glyph, i.e. the glyph colors are always the same regardless of the current fill or stroke color. OpenType color fonts with a COLR table may contain glyph layers which are designated as »foregroundcolor« instead of using intrinsic color. This color and the corresponding transparency can be supplied with the *foregroundcolor* and *foregroundtransparency* font options; these values are fixed for all glyphs in the font and all uses of such glyphs.

Color glyphs can be skipped with `colormode=ignorecolor`; in this case the font is processed like a font which doesn't contain any color information.

- ▶ Monochromatic (or mono) glyphs which describe only the shape, but don't contain any color or transparency information. Mono glyphs can be colorized with the current fill color when they are placed on the page. Mono glyphs can be skipped with `colormode=ignoremono`. However, keep in mind that not all color fonts contain mono glyphs. Most *sbix* color fonts don't contain usable mono glyphs and are therefore not suitable for use with `colormode=ignorecolor`.

Color and mono versions of a particular glyph may be present in a font simultaneously.

Using color fonts. Some notes on using OpenType color fonts in applications:

- ▶ Many emoji characters are encoded outside the BMP, i.e. at Unicode code points > U+FFFF. Depending on your development environment different methods are available for expressing such Unicode values. More information and examples can be found in the language-specific sections of Chapter 2.
- ▶ Emoji characters can be modified and combined in various ways, see Section 7.3, »Emoji Sequences«, page 138.
- ▶ OpenType color fonts with emoji characters benefit from the fallback font feature; see »Add color emoji characters to a text font«, page 128 for an example.

- ▶ Unlike other font formats, subsetting is enabled by default for color fonts. Subsetting can be disabled with the option `subsetting=false`.
- ▶ Since color glyphs from OpenType color fonts have intrinsic color and transparency, these aspects cannot be modified when using the font. Similarly, the text appearance option `textrendering` doesn't have any effect on color glyphs except that `textrendering=3` (invisible text) and 7 (clipping) result in invisible text.
- ▶ OpenType color fonts may contain glyphs with transparent elements. This is relevant for PDF/A-1 and PDF/X-3 where transparency is not allowed. If subsetting is enabled, transparent glyphs are not detected when used, but only in `PDF_end_document()` when the font subset is created. Because of this the page option `transparencysgroup=auto` doesn't detect transparency in color fonts when subsetting is enabled.
- ▶ Unicode variation sequences can be used to select the text or emoji style glyph for a particular character (see Section 7.7, »Unicode Variation Sequences (UVS)«, page 153).

Querying information about color fonts. After successfully loading an OpenType color font the following color-related information can be retrieved with `PDF_info_font()`:

- ▶ The keyword `colortype` retrieves the type of an OpenType color font.
- ▶ The keyword `colormode` retrieves the processing mode.
- ▶ The keyword `numcolorglyphs` retrieves the number of color glyphs.

6.1.4 WOFF and WOFF2 Fonts

WOFF (Web Open Font Format) is a compressed container format for TrueType and OpenType fonts. WOFF has been designed for use on the Web and offers compression and subsetting features to achieve small font file sizes. WOFF fonts use the file name extension `.woff`.



The WOFF2 format is based on WOFF and provides improved compression. WOFF2 has been published as W3C recommendation in 2018. WOFF2 fonts use the file name extension `.woff2`.

PDFlib supports WOFF and WOFF2 fonts provided the underlying TrueType or OpenType font is supported.

Note WOFF2 fonts are not supported on IBM System i and IBM Z.

6.1.5 PostScript Type 1 Fonts (deprecated)

Note Support for PostScript Type 1 fonts is deprecated.

PostScript Type 1 fonts are split in two parts: outline data and metrics information. PDFlib supports the following formats for PostScript Type 1 fonts:

- ▶ The platform-independent AFM (Adobe Font Metrics) and the Windows-specific PFM (Printer Font Metrics) format for metrics information.
- ▶ The platform-independent PFA (Printer Font ASCII) and the Windows-specific PFB (Printer Font Binary) format for outline information in Type 1 format.

6.1.6 SING Fonts (Glyphlets)

SING fonts (*Smart Independent Glyphlets*) are an extension of the OpenType font format. SING fonts have been developed as a solution to the Gaiji problem with CJK text, i.e. custom glyphs which are not encoded in Unicode.

SING fonts contain only a single glyph (they may also contain an additional vertical variant). The Unicode value of this »main« glyph can be retrieved with PDFlib by requesting its glyph ID and subsequently the Unicode value for this glyph ID:

```
maingid = (int) p.info_font(font, "maingid", "");
uv = (int) p.info_font(font, "unicode", "gid=" + maingid);
```

It is recommended to use SING fonts as fallback font with the *gaiji* suboption of the *forcechars* option of the *fallbackfonts* option of *PDF_load_font()*; see Section 7.6.3, »EUDC and SING Fonts for Gaiji Characters«, page 151.

6.1.7 Type 3 Fonts

Unlike all other font formats, Type 3 fonts are not fetched from a file, but are defined at runtime with PDFlib graphics methods. Type 3 fonts are useful for the following purposes:

- ▶ bitmap fonts;
- ▶ fonts with intrinsic color or transparency;
- ▶ custom graphics such as logos can easily be printed using simple text operators;
- ▶ Japanese gaiji (user-defined characters) which are not available in any font or encoding.
- ▶ Type 3 fonts are created internally for color fonts (see Section 6.1.3, »OpenType Color Fonts in SVG, COLR and sbix Flavors«, page 110).

Type 3 fonts may contain two types of glyphs:

- ▶ By default, glyphs are uncolorized: the glyph description specifies only the shape of the glyph. The color is specified with the current fill color when text in the generated font is emitted. Methods for setting the color or transparency as well as images (except masks) are not allowed in the glyph description.
- ▶ If the glyph option *colorized* is set to *true*, glyph descriptions specify color and transparency and may also place arbitrary images. The colors are fixed for the glyph, i.e. the glyph will appear identically regardless of the fill color in effect when the text is placed.

Since all PDFlib features for vector graphics, raster images, and even text output can be used in Type 3 font definitions (subject to the color restrictions described above), there are no restrictions regarding the contents of glyphs in a Type 3 font. Type 3 fonts are most often used for bitmapped glyphs. The following example demonstrates the definition of a simple Type 3 font:

```
p.begin_font("Fuzzyfont", 0.001, 0.0, 0.0, 0.001, 0.0, 0.0, "");

p.begin_glyph_ext(-1, "glyphname=circle width=1000 boundingbox={0 0 1000 1000}");
p.arc(500, 500, 500, 0, 360);
p.fill();
p.end_glyph();

p.begin_glyph_ext(-1, "glyphname=ring width=400 boundingbox={0 0 400 400}");
p.arc(200, 200, 200, 0, 360);
p.stroke();
p.end_glyph();

p.end_font();
```

Cookbook Full code samples can be found in the *Cookbook topics* `type3_fonts/starter_type3font`, `type3_fonts/type3_bitmaptext`, and `type3_fonts/type3_vectorlogo`.

The font is registered in PDFlib and its name can be supplied to `PDF_load_font()`. Note the following when working with Type 3 fonts:

- ▶ The `.notdef` glyph must be present and must always occupy slot 0.
- ▶ It is strongly recommended to supply proper Unicode values to `PDF_begin_glyph_ext()`. PDFlib automatically creates correct glyph names according to the AGL (Adobe Glyph List) from the supplied Unicode values, but alternate names can be supplied in the `glyphname` option.
- ▶ It is recommended to use the `inline` image option (only for `imagetype=jpeg` and `raw`) for defining glyphs with bitmaps in Type 3 fonts. The `interpolate` option for images may be useful for enhancing the appearance of Type 3 bitmap fonts.
- ▶ When bitmap data is used to define glyphs, unused pixels in the bitmap print as white, regardless of the background. In order to avoid this and have the background color shine through, use the `mask` option for constructing the bitmap image.
- ▶ Since Type 3 fonts in PDF require 8-bit addressing they are limited to 255 usable glyphs (since the `.notdef` glyph is required). However, PDFlib accepts an arbitrary number of glyphs for a Type 3 font and creates as many auxiliary Type 3 fonts as required for representing the whole glyph set. Type 3 fonts loaded with `unicode` encoding are not subject to the 255 glyph restriction.

6.2 Unicode Characters and Glyphs

6.2.1 Glyph IDs

A font is a collection of glyphs, where each glyph defines a visible shape. PDFlib assigns a number to each glyph in the font. This number is called the glyph id or GID. GID 0 (zero) refers to the `.notdef` glyph. The visual appearance of the `.notdef` glyph varies among font formats and vendors; typical implementations are the space glyph or a hollow or crossed-out rectangle. The highest GID is one less than the number of glyphs in the font which can be queried with the `numglyphs` keyword of `PDF_info_font()`.

The assignment of glyph IDs depends on the font format:

- ▶ Since TrueType and OpenType fonts already contain internal GIDs, PDFlib uses these GIDs.
- ▶ For CID-keyed OpenType CJK fonts CIDs are used as GIDs.
- ▶ For other font types PDFlib numbers the glyphs according to the order of the corresponding outline descriptions in the font.

PDFlib supports glyph selection via GID as an alternative to Unicode and other encodings (see »Glyphid encoding«, page 117). Direct GID addressing is useful for specialized applications, e.g. printing font overview tables by querying the number of glyphs and iterating over all glyph IDs.

6.2.2 Unicode Mapping for Glyphs

Unicode mapping and ambiguous glyphs. PDFlib assigns Unicode values to all glyphs in a font. Unicode values may not be available for all glyphs in the font, e.g. for variant

glyphs (e.g. swash characters), extended ligatures, and non-textual symbols outside the Unicode standard. In this case PDFlib assigns PUA values to the affected glyphs.

In some fonts a particular glyph may be used to represent multiple Unicode values. Common examples for such ambiguous glyphs are the empty glyph which represents U+0020 Space as well as U+00A0 No-Break Space, or a glyph which represents both U+2126 Ohm Sign and U+03A9 Greek Capital Letter Omega.

If multiple Unicode values are represented by the same glyph PDFlib creates a ToUnicode CMap which maps the glyph to one of the affected Unicode values. The other Unicode values are emitted with the same glyph ID, but are assigned an *ActualText* attribute with the appropriate Unicode value. This way correct semantics are preserved in the generated PDF output.

Text selection problems in Acrobat. Unfortunately, Acrobat DC sometimes has problems with the text selection of text with an *ActualText* attribute. Some glyphs cannot be selected and highlighted. Although the visual highlight is incomplete, the text contained in the *ActualText* attribute is copied to the clipboard correctly. You can work around the selection problems in the following ways:

- ▶ Avoid ambiguous glyphs in the first place.
- ▶ Use the text filter option *actualtext=false* to disable creation of the *ActualText* attribute.

However, these methods are not recommended since different characters may be extracted with the same Unicode value which thwarts content repurposing.

Unmapped glyphs and the Private Use Area (PUA). In some situations the font may not provide a Unicode value for a particular glyph. In this case PDFlib assigns a value from the Unicode Private Use Area (PUA, see Section 5.1, »Important Unicode Concepts«, page 95) to the glyph. Such glyphs are called *unmapped glyphs*. The number of unmapped glyphs in a font can be queried with the *unmappedglyphs* keyword of *PDF.info_font()*. Unmapped glyphs are represented by the Unicode replacement character U+FFFD in the font's ToUnicode CMap which controls searchability and text extraction. As a consequence, unmapped glyphs cannot be properly extracted as text from the generated PDF. However, this behavior can be changed which is particularly useful for CJK Gaiji characters; see »Preserving PUA values for Gaiji characters«, page 151, for details.

When PDFlib assigns PUA values to unmapped glyphs it uses ascending values from the following pool:

- ▶ The basis is the Unicode PUA range in the Basic Multilingual Plane (BMP), i.e. the range U+E000 - U+F8FF. Additional PUA values in plane 15 (U+F0000 to U+FFFFFD) are used if required.
- ▶ PUA values which have already been assigned by the font internally are not used when creating new PUA values.
- ▶ PUA values in the Adobe range U+F600-F8FF are not used.

The generated PUA values are unique within a font. The assignment of generated PUA values for the glyphs in a font is independent from other fonts.

6.3 Loading Fonts

6.3.1 Selecting an Encoding for Text Fonts

Fonts can be loaded explicitly with the `PDF_load_font()` method or implicitly by supplying the `fontname` and `encoding` options to methods such as `PDF_add/create_textflow()` or `PDF_fill_textblock()`. Regardless of the method used for loading a font a suitable encoding must be specified. The encoding determines

- ▶ in which text formats PDFlib expects the supplied text;
- ▶ which glyphs in a font can be used;
- ▶ how text on the page and the glyph data in the font is stored in the PDF output.

PDFlib's text handling is based on the Unicode standard. Since all modern development environments support Unicode our goal is to make it as easy as possible to use Unicode strings for creating PDF output. However, developers who don't work with Unicode are not required to switch their application to Unicode since legacy encodings can be used as well.

The choice of encoding depends on the font, the available text data and some programming aspects. In the remainder of this section we provide an overview of the different classes of encodings as an aid for selecting a suitable encoding.

Unicode encoding. With `encoding=unicode` you can pass Unicode strings to PDFlib. This encoding is supported for all font formats. Depending on the language binding in use, the Unicode string data type provided by the programming language (e.g. Java) can be used, or byte arrays containing Unicode in one of the UTF-8, UTF-16, or UTF-32 formats with little- or big-endian byte ordering (e.g. C).

With `encoding=unicode` all glyphs in a font can be addressed; complex script shaping and OpenType layout features are supported. PDFlib checks whether the font contains a glyph for a requested Unicode value. If no glyph is available, a substitute glyph can be pulled from another font (see Section 6.3.6, »Fallback Fonts«, page 127).

In non-Unicode-capable language bindings PDFlib expects UTF-16 encoded text by default. However, you can supply single-byte strings by specifying `textformat=bytes`. In this case the byte values represent the characters U+0001 - U+00FF, i.e. the first Unicode block with Basic Latin characters (identical to ISO 8859-1). Using character references Unicode values outside this range can be specified in single-byte text.

Single-byte encodings. 8-bit encodings map each byte in a text string to a single character and are thus limited to 255 different characters at a time (the value 0 is not available). This type of encoding is supported for all font formats. PDFlib checks whether the font contains glyphs which match the selected encoding. If a minimum number of usable glyphs is not reached, PDFlib logs a warning message. If no usable glyph at all for the selected encoding is available in the font, font loading fails with the message *font doesn't support encoding*. PDFlib checks whether the font contains a glyph for a requested input value. If no glyph is available, a substitute glyph can be pulled from another font (see Section 6.3.6, »Fallback Fonts«, page 127).

In non-Unicode-capable language bindings PDFlib expects single-byte encoded text by default. However, you can supply UTF-8 or UTF-16 strings by specifying `textformat=utf8` or `utf16`.

8-bit encodings are discussed in detail in Section 5.3.3, »Single-Byte (8-Bit) Encodings«, page 101. They can be pulled from various sources:

- ▶ A large number of predefined encodings. These cover the most important encodings in use on a variety of systems and in a variety of locales.
- ▶ User-defined encodings which can be supplied in an external file or constructed dynamically at runtime with *PDF_encoding_set_char()*. These encodings can be based on glyph names or Unicode values.
- ▶ Encodings pulled from the operating system, also known as a *system encoding*. This feature is available on Windows, IBM System i and IBM Z.

Builtin encoding. You can specify *encoding=builtin* to use single-byte codes for non-textual glyphs from symbolic fonts. A font's builtin encoding depends on the font type:

- ▶ OpenType font with TrueType outlines: the encoding is created based on the font's symbolic cmap, i.e. the (3, 0) entry in the cmap table. If this entry is not present *encoding=builtin* cannot be used.
- ▶ OpenType CFF fonts can contain an encoding in the CFF table.
- ▶ For Type 3 fonts the builtin encoding is defined by the first 255 glyphs in the font.

If the font does not contain any builtin encoding, font loading fails (e.g. OpenType CJK fonts). You can use the *symbolfont* key in *PDF_info_font()*. If it returns *false*, the font is a text font which can be loaded with one of the common encodings. This is not possible if the *symbolfont* key returns *true*. The glyphs in symbolic fonts can only be used if you know the corresponding code for each glyph (see Section 6.3.2, »Selecting an Encoding for symbolic Fonts«, page 117).

In non-Unicode-capable language bindings PDFlib expects single-byte text by default. This has the advantage that you can use the single-byte values which have traditionally been used to address some symbolic fonts; this is not possible with other encodings. However, you can also supply text in a Unicode format, e.g. with *textformat=utf16*.

The disadvantage of *encoding=builtin* is that in single-byte encoded text character references cannot be used.

Glyphid encoding. PDFlib supports *encoding=glyphid* for all font formats. With this encoding all glyphs in a font can be addressed with the numbering scheme explained in Section 6.2.1, »Glyph IDs«, page 114. Numerical glyph IDs may run from 0 to 65 565. The maximum glyph ID value can be queried with the *maxcode* key in *PDF_info_font()*.

In non-Unicode-capable language bindings PDFlib expects double-byte encoded text by default (*textformat=utf16*).

PDFlib checks whether the supplied glyph ID is valid for the font. Complex script shaping and OpenType layout features are supported.

Since glyph IDs are font-specific and in some situations are even created by PDFlib, *encoding=glyphid* is generally not suited for regular text output. The main use of this encoding is for applications which already work with glyph IDs and for printing font tables with all glyphs.

6.3.2 Selecting an Encoding for symbolic Fonts

Symbolic fonts are fonts which contain symbols, logos, pictograms or other non-textual glyphs. They raise several issues which are not relevant for text fonts. The underlying problem is that by design the Unicode standard does not generally encode symbolic

glyphs (although there are exceptions to this rule, e.g. the glyphs in the ZapfDingbats font). In order to make symbolic fonts fit for use in Unicode workflows, TrueType and OpenType fonts usually assign Unicode values in the Private Use Area (PUA) to their glyphs. For lack of Unicode mapping tables, PostScript Type 1 fonts cannot do this, and generally use the codes of Latin characters to select their glyphs. In all font formats the symbolic glyphs usually have custom glyph names.

This situation has the following consequences regarding selection of glyphs from symbolic fonts:

- ▶ Symbolic TrueType and OpenType fonts are best loaded with *encoding=unicode*. If you know the PUA values assigned to the glyphs you can supply these values in the text in order to select symbolic glyphs. This requires advance knowledge of the PUA assignments in the font.
- ▶ Since PDFlib assigns PUA values for symbolic PostScript Type 1 fonts internally, these PUA values are not known in advance.
- ▶ If you prefer to work with 8-bit codes for addressing the glyphs in a symbolic font you can load the font with *encoding=builtin* and supply the 8-bit codes in the text. For example, the digit 4 (code 0x34) will select the check mark symbol in the ZapfDingbats font.

In order to use symbolic fonts with *encoding=unicode* suitable Unicode values must be used for the text:

- ▶ Microsoft's symbolic fonts, e.g. Wingdings and Webdings, use PUA Unicode values in the range U+F020 - U+F0FF.
- ▶ For other fonts the Unicode values for individual glyphs in the font must be known in advance or must be determined at runtime with *PDF_info_font()*.

Control characters. The Unicode control characters in the range U+0001 - U+001F which are listed in Table 9.1 are supported in Textflow even with *encoding=builtin*. Codes < 0x20 are interpreted as control characters if the symbolic font does not contain any glyph for the code. This is true for the majority of symbolic fonts.

Since the code for the a linefeed characters differs between ASCII and EBCDIC, it is recommended to avoid the literal character 0x0A on EBCDIC systems and use the PDFlib escape sequence `\n` with the option *escapesequence=true* instead. Note that the `\n` must arrive at the PDFlib API, e.g. in C the sequence `\\n` is required.

Character references. Character references are supported for symbolic fonts. However, symbolic fonts generally do not include any glyph for the ampersand character U+0026 '&' which introduces character references. The code 0x26 cannot be used either since it could be mapped to an existing glyph in the font. For these reasons symbolic fonts should be loaded with *encoding=unicode* if character references must be used. Character references do not work with *encoding=builtin*.

6.3.3 Example: Selecting a Glyph from the Wingdings Symbol Font

Since there are many different ways of selecting characters from a symbol font and some will not result in the desired output, let's take a look at an example.

Understanding the characters in the font. First let's collect some information about the target character in the font, using the Windows *charmap* application (see figure 6.1):

- ▶ *Charmap* displays the glyphs in the Wingdings font, but does not provide any Unicode access in the *Advanced view*. This is a result of the fact that the font contains symbolic glyphs for which no standardized Unicode values are registered. Instead, the glyphs in the font use dummy Unicode values in the Private Use Area (PUA). The *charmap* application does not reveal these values.

- ▶ If you look at the lower left corner of the *charmap* window or hover the mouse over the *smileface* character, the *Character code: 0x4A* is displayed. This is the glyph's byte code.

This code corresponds to the uppercase *J* character in WinAnsi encoding. For example, if you copy the character to the clipboard the corresponding Unicode value U+004A, i.e. character *J* results from pasting the clipboard contents. Nevertheless, this is not the character's Unicode value and therefore U+004A or *J* cannot be used to select it in Unicode workflows.

- ▶ The Unicode character used in the font is not displayed in *charmap*. However, symbolic fonts provided by Microsoft use the following simple rule:

Unicode value = U+F000 + (character code displayed in *charmap*)

For the *smileface* glyph this yields the Unicode value U+F04A.

- ▶ The corresponding glyph name can be retrieved with a font editor and similar tools. In our example it is *smileface*.

You can use *PDF_info_font()* to query Unicode values, glyph names or codes, see Section 6.5.2, »Font-specific Encoding, Unicode, and Glyph Name Queries«, page 132.

Addressing the symbol character with PDFlib. Depending on the information which is available about the target character you can select the Wingdings *smileface* glyph in several ways:

- ▶ If you know the PUA Unicode value assigned to the character in the font you can use a numerical character reference (see »Numerical character references«, page 106):

```
&#xF04A;
```

If you work with UTF-8 you can use the corresponding three-byte sequence:

```
\xEF\x81\x8A
```

Unicode values cannot be used with the combination of *encoding=builtin* and *textformat=bytes*.

- ▶ If you know the character code you can use a byte value reference (see »Byte value references«, page 107):

```
&.#x4A;
```

In non-Unicode-capable language bindings the character code can be specified directly if *encoding=builtin* and *textformat=bytes*:

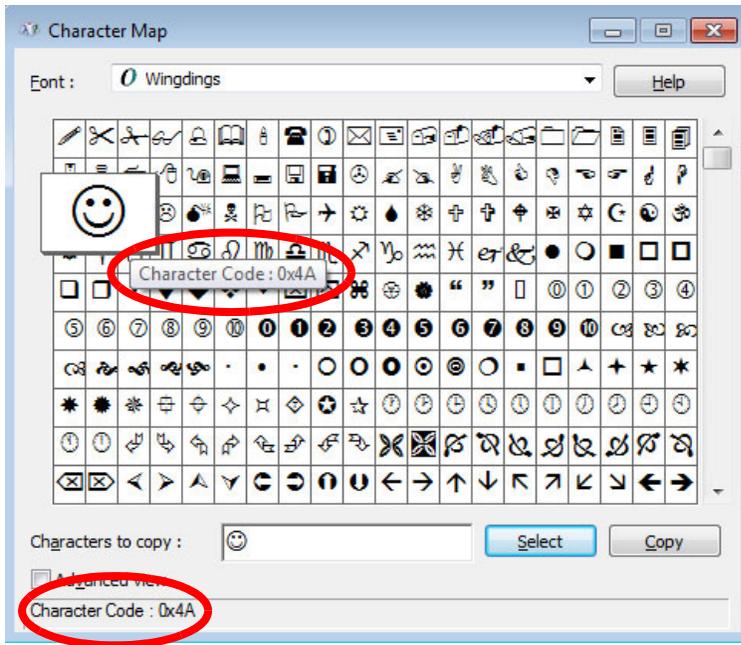


Fig. 6.1
Windows character map
with the Wingdings font

J
\x4A

- ▶ If you know the glyph name you can use a glyph name reference (see »Glyph name references«, page 107):

&.smileface;

Glyph names cannot be used with the combination of *encoding=builtin* and *text-format=bytes*.

Table 6.2 lists methods for Unicode-capable language bindings such as Java and .NET.

Table 6.2 Addressing the smileface glyph in the Wingdings font with Unicode-capable language bindings (e.g. Java)

encoding	additional options	input string	visible result on the page
unicode		<code>\uF04A</code> ¹	☺
	charref	<code>&#xF04A;</code>	☺
	charref	<code>&#x4A;</code>	☺
	charref	<code>&.smileface;</code>	☺
		<code>J</code> ²	(space or .notdef)
builtin		<code>\x4A</code>	(space or .notdef)
	(same as above with <i>encoding=unicode</i>)		

1. String syntax for U+F04A in Java and many other Unicode-capable languages

2. Winansi character for the byte code \x4A

Table 6.3 lists methods for non-Unicode-capable language bindings such as C.

Table 6.3 Addressing the smileface glyph in the Wingdings font with non-Unicode-capable language bindings, e.g. C

encoding	textformat	additional options	input string	visible result on the page	
unicode	utf16		<code>\xF0\x4A</code> ¹	☺	
	utf8, bytes	charref	<code>&#xF04A;</code>		☺
		charref	<code>&#x4A;</code>		☺
		charref	<code>&.smileface;</code>		☺
				<code>if textformat=utf8: i•Š</code> ²	☺
				<code>if textformat=bytes: J</code> ³	(space or .notdef)
		escapesequence ⁴		<code>if textformat=utf8: \xEF\x81\x8A</code> ⁵	☺
escapesequence		<code>if textformat=bytes: \x4A</code>	(space or .notdef)		
builtin	utf16, utf8	(same as above with encoding=unicode)			
	bytes	charref	<code>&#xF04A;</code>		
		charref	<code>&#x4A;</code>		
		charref	<code>&.smileface;</code>		
			<code>J</code>	☺	
		escapesequence	<code>\x4A</code>	☺	

1. Must be expressed as `\xF0\x4A` or `\x4A\xF0` depending on byte ordering; note that `\x` designates C escape syntax
2. Winansi characters for the three-byte sequence `\xEF \x81 \x8A`
3. Winansi character for the byte code `\x4A`
4. The escapesequence option is only required if the programming language doesn't offer any syntax for direct byte values.
5. Three-byte UTF-8 sequence for U+F04A

6.3.4 Searching for Fonts

Sources of Font Data. As mentioned earlier, fonts can be loaded explicitly with the `PDF_load_font()` method or implicitly by supplying the `fontname` option to various text output functions. You can use a font's native name or work with arbitrary custom names which will be used to locate the font data. Custom font names must be unique within a document. In `PDF_info_font()` this font name can be queried with the `apiname` key.

Subsequent calls to `PDF_load_font()` with the same font name will return the same font handle if all options are identical to those provided in the first call to this method (a few options are treated differently; see PDFlib API Reference for details). Otherwise a new font handle is created for the same font name. PDFlib supports the following sources of font data:

- ▶ Disk-based or virtual font files
- ▶ Fonts pulled from the Windows or macOS operating system (host fonts)
- ▶ PDF Standard fonts: these are from a small set of Latin fonts with well-known names
- ▶ Type 3 fonts which have been defined with `PDF_begin_font()` and related functions.

Cookbook A full code sample can be found in the *Cookbook* topic `fonts/font_configuration`.

Automatically creating a font resource list. With the `enumeratefonts` option PDFlib can be instructed to collect all fonts which are accessible on the search path (see »File search and the SearchPath resource category«, page 52). Using the `saveresources` option the current list of PDFlib resources can be written to a file:

```
/* Add font directory to the search path */
p.set_option("searchpath={{C:/fonts} {D:/fonts}}");

/* Enumerate all fonts on the searchpath and create a UPR file */
p.set_option("enumeratefonts saveresources={filename=C:/pdflib.upr}");
```

As a result all fonts available in the specified directories are listed in the generated UPR file for later use.

Font name aliasing. Each font may have an arbitrary number of alias names. This may be useful in situations where fonts are requested via an artificial or virtual name which must be mapped to a physical font. Font name aliases can be created with the `FontnameAlias` resource category (see Table 3.1, page 51) as in the following example:

```
p.set_option("FontnameAlias={sans Helvetica}");
```

The alias name to the left can be chosen arbitrarily and can be used for loading the font under its new alias name. The name on the right must be a valid API name of a font, e.g. the name of a host font or a font which has been connected to a font resource with one of the font resource categories `FontOutline` etc.

Search order for fonts. The font name supplied to PDFlib is a name string. If the specified name is a font name alias it will be replaced with the corresponding API font name. PDFlib uses the API font name to search for fonts of various types in the order described below. The search process stops as soon as one of the steps located a usable font:

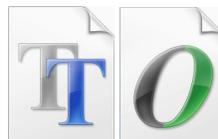
- ▶ The font name matches the name of a Type 3 font which has been created with *PDF_begin_font()* (see Section 6.1.7, »Type 3 Fonts«, page 113).
- ▶ The font name matches the name in a *FontOutline* resource which connects the font name with the name of a TrueType or OpenType font file.
- ▶ The font name matches the name in a *HostFont* resource which connects the font name with the name of a font installed on the system.
- ▶ The font name matches the name of a builtin CJK font (see »Builtin CJK Fonts«, page 124) and *embedding=false*.
- ▶ The font name matches the name of a Latin core font (see »Standard Type 1 fonts (Latin core fonts)«, page 124) and *embedding=false*.
- ▶ The name matches the name of a host font installed on the system (see Section 6.3.5, »Host Fonts on Windows and macOS«, page 126).
- ▶ The font name matches the base name (i.e. without file name suffix) of a font file.

If no font was found, font loading stops with the message *Font not found*.

Details regarding the resource categories can be found in Section 3.1.4, »Resource Configuration and File Search«, page 51. The following sections discuss font loading for the various classes of fonts in more detail.

TrueType, OpenType, and WOFF/WOFF2 fonts. The font name must be connected to the name of the desired font file via the *FontOutline* resource:

```
p.set_option("FontOutline={Arial=/usr/fonts/Arial.ttf}");
font = p.load_font("Arial", "unicode", "");
```



The font name to the left of the equal sign (called the font's API name) can be chosen arbitrarily:

```
p.set_option("FontOutline={f1=/usr/fonts/Arial.ttf}");
font = p.load_font("f1", "unicode", "");
```



As an alternative to runtime configuration via *PDF_set_option()*, the *FontOutline* resource can be configured in a UPR file (see Section 3.1.4, »Resource Configuration and File Search«, page 51). In order to avoid absolute file names you can use the *SearchPath* resource category (again, the *SearchPath* resource category can alternatively be configured in a UPR file), for example:

```
p.set_option("SearchPath={{/usr/fonts}}");
p.set_option("FontOutline={f1=Arial.ttf}");
font = p.load_font("f1", "unicode", "");
```

TrueType and OpenType Collections. In order to select a font which is contained in a TrueType or OpenType Collection (TTC, see Section 7.6.1, »Using TrueType and OpenType CJK Fonts«, page 150) file you directly specify the name of the font:



```
p.set_option("FontOutline={MS-Gothic=msgothic.ttc}");
font = p.load_font("MS-Gothic", "unicode", "");
```

The font name is matched against the names of all fonts in the TTC/OTC file. Alternatively, to select the *n*-th font in a TTC/OTC file you can specify the number *n* with a co-

lon after the font name. In this case the API font name to the left of the equal sign can be chosen arbitrarily:

```
p.set_option("FontOutline={f1=msgothic.ttc}");
font = p.load_font("f1:0", "unicode", "");
```

Standard Type 1 fonts (Latin core fonts). PDFlib supports a set of 14 standard fonts which don't have to be configured since metrics information for these fonts is built into PDFlib. Corresponding font files are only required if the fonts are to be embedded. The core fonts have the following names:

Courier, Courier-Bold, Courier-Oblique, Courier-BoldOblique, Helvetica, Helvetica-Bold, Helvetica-Oblique, Helvetica-BoldOblique, Times-Roman, Times-Bold, Times-Italic, Times-BoldItalic, Symbol, ZapfDingbats

If a font name is not connected to any file name via resources, PDFlib searches the font in the list of Standard Type 1 fonts.

The Standard Type 1 fonts can only be used with *embedding=false*. This implies that the fonts cannot be used for creating PDF/A output which requires embedding.

Builtin CJK Fonts. Similar to the Standard Type 1 fonts which cover the Latin script, PDFlib includes internal font metrics for a set of CID fonts which cover the Chinese, Japanese and Korean scripts in horizontal and vertical writing mode. These fonts can be used without external fonts files provided the conditions below are observed. They can be used with *encoding=unicode* or (deprecated) the name of a predefined CMap for legacy encodings. The names of the builtin CJK fonts are as follows:

Simplified Chinese: *AdobeSongStd-Light*

Traditional Chinese: *AdobeMingStd-Light*

Japanese: *KozGoPro-Medium* (sans serif) and *KozMinProVI-Regular* (serif)

Korean: *AdobeMyungjoStd-Medium*

The following requirements must be met to use these fonts:

- ▶ The font option *embedding=false* must be supplied. This implies that the fonts cannot be used for creating PDF/A output which requires embedding.
- ▶ (Deprecated) The name of a predefined CMap for legacy encodings can be used as encoding.
- ▶ For *encoding=unicode* the CMap *Adobe-*-UCS2* plus the following CMaps for the applicable character collection GB1, CNS1, Japan1, or Korea1 must be available:
compatibility=1.4: Uni-UCS2-H* or (with *vertical=true*) *Uni*-UCS2-V*
compatibility > 1.4: Uni-UTF16-H* or (with *vertical=true*) *Uni*-UTF16-V*

For example, *KozGoPro-Medium* for horizontal writing mode requires the following CMaps:

- ▶ *compatibility > 1.4: Adobe-Japan1-UCS2* and *UniJIS-UTF16-H*
- ▶ *compatibility=1.4: Adobe-Japan1-UCS2* and *UniJIS-UCS2-H*

Host fonts. If a font name is not connected to any file name via resources, PDFlib searches the font in the list of fonts installed on Windows or macOS. Fonts installed on the system are called *host fonts*. Host font names must be encoded in ASCII. On Windows

Unicode can also be used. See Section 6.3.5, »Host Fonts on Windows and macOS«, page 126, for more details on host fonts. Example:

```
font = p.load_font("Verdana", "unicode", "");
```

On Windows an optional font style can be added to the font name after a comma (this syntax can also be used with the Latin core fonts):

```
font = p.load_font("Verdana,Bold", "unicode", "");
```

In order to load a host font with the name of one of the core fonts, the font name must be connected to the desired host font name via the *HostFont* resource category. The following fragment makes sure that instead of using the built-in core font data, the Symbol font data is taken from the host system:

```
p.set_option("HostFont={Symbol=Symbol}");  
font = p.load_font("Symbol", "unicode", "");
```

Extension-based search for font files. All font types except Type 3 fonts can be searched by using the specified font name as the base name (without any suffix) of a font file. If PDFlib cannot find any font with the specified name it will loop over all directories listed in the *SearchPath* resource category and add all known file name suffixes to the supplied font name in an attempt to locate the font file. The details of the extension-based search algorithm are as follows:

- ▶ The following suffixes are added to the font name, and the resulting file names tried one after the other to locate the font file:

```
.ttf .otf .woff .ttf.woff .otf.woff .woff2 .ttf.woff2 .otf.woff2 .ttc .otc .tte .gai  
.TTF .OTF .WOFF .TTF.WOFF .OTF.WOFF .WOFF2 .TTF.WOFF2 .OTF.WOFF2 .TTC .OTC .TTE .GAI
```

- ▶ All candidate file names above are searched for »as is«, and then by prepending all directory names configured in the *SearchPath* resource category.

This means that PDFlib will find a font without any manual configuration provided the corresponding font file name consists of the font name plus the standard file name suffix according to the font type and is located in one of the *SearchPath* directories.

The following groups of statements achieve the same effect with respect to locating the font outline file:

```
p.set_option("FontOutline={Arial=/usr/fonts/Arial.ttf}");  
font = p.load_font("Arial", "unicode", "");
```

and

```
p.set_option("SearchPath={{/usr/fonts}}");  
font = p.load_font("Arial", "unicode", "");
```

Type 3 fonts. Type 3 fonts must be defined at runtime by defining the glyphs with PDFlib graphics methods (see Section 6.1.7, »Type 3 Fonts«, page 113). If the font name supplied to *PDF_begin_font()* matches the font name requested with *PDF_load_font()* the font is selected at the beginning of the font search. Example:

```
p.begin_font("PDFlibLogoFont", 0.001, 0.0, 0.0, 0.001, 0.0, 0.0, "");
```

```
...
```

```
p.end_font();
...
font = p.load_font("PDFlibLogoFont", "logoencoding", "");
```

6.3.5 Host Fonts on Windows and macOS

On macOS and Windows PDFlib can access fonts which have been installed in the operating system. We refer to such fonts as *host fonts*. Instead of manually configuring font files simply install the font in the system and PDFlib will use it.

When working with host fonts it is important to use the exact case-sensitive font name. Host font search can be disabled with the option `usehostfonts=false` of `PDF_set_option()`.

Host font names on macOS. the standard TrueType/OpenType font name can be used to retrieve the font, i.e. the »full font name« in the font's name table. You can use the FontBook application to determine the names of installed fonts.

Host font names on Windows. You can find the name of an installed font by double-clicking the font file and taking note of the full font name which is displayed in the window title. Some fonts may have parts of their name localized according to the respective Windows version in use. For example, the common font name portion *Bold* may appear as the translated word *Fett* on a German system. In order to retrieve the host font data from the Windows system you must use the translated form of the font name in PDFlib (e.g. *Arial Fett*), or use font style names (see below). However, in order to retrieve the font data directly from file you must use the generic (non-localized) form of the font name (e.g. *Arial Bold*).

You can avoid this internationalization problem by appending font style names (e.g. »,Bold«, see below) to the font name instead of using localized font name variants.

Windows font style names. When loading host fonts from the Windows operating system PDFlib users have access to a feature provided by the Windows font selection machinery: style names can be provided for the weight and slant, for example

```
font = p.load_font("Verdana,Bold", "unicode", "");
```

This instructs Windows to search for a particular bold, italic, or other variation of the base font. Depending on the available fonts Windows selects a font which most closely resembles the requested style (it does not create a new font variation). The font found by Windows may be different from the requested font, and the font name in the generated PDF may be different from the requested name; PDFlib does not have any control over Windows' font selection. Font style names only work with host fonts, but not for fonts configured via a font file.

The following keywords (separated from the font name with a comma) can be attached to the base font name to specify the font weight:

```
none, thin, extralight, ultralight, light, normal, regular, medium,
semibold, demibold, bold, extrabold, ultrabold, heavy, black
```

The keywords are case-insensitive. The *italic* keyword can be specified alternatively or in addition to the above. If two style names are specified each must be separated with a comma, for example:

```
font = p.load_font("Verdana,Bold,Italic", "unicode", "");
```

Numerical font weight values are equivalent to font style names:

0 (none), 100 (thin), 200 (extralight), 300 (light), 400 (normal), 500 (medium), 600 (semibold), 700 (bold), 800 (extrabold), 900 (black)

The following example selects the bold variant of a font:

```
font = p.load_font("Verdana,700", "unicode", "");
```

Note Windows style names for fonts may be useful if you have to deal with localized font names since they select font styles regardless of their localized names.

Windows font substitution. Windows can automatically substitute fonts based on certain registry entries. This kind of font substitution also affects PDFlib's host font mechanism, but is completely under control of the Windows operating system. For example, if the Helvetica font is requested Windows may deliver Arial instead, depending on the following registry entry:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\FontSubstitutes
```

For more information about Windows font substitution refer to Microsoft documentation.

6.3.6 Fallback Fonts

Note Fallback fonts are not supported for `encoding=glyphid`.

Cookbook A full code sample can be found in the Cookbook topic `fonts/starter_fallback`.

Fallback fonts provide a powerful mechanism which deals with shortcomings in fonts and encodings. They are also useful for combining fonts for different scripts. For example, a font for the Arabic script may not support Latin characters. Fallback fonts can be used in many situations to facilitate text output since necessary font switching is done automatically by PDFlib. This mechanism augments a given font (called the base font) by using glyphs from one or more other fonts in addition to glyphs from the base font. Fallback fonts offer the following features:

- ▶ Glyphs which are unavailable in the base font are automatically searched in one or more fallback fonts. In other words, you can add glyphs to a font. Since multiple fallback fonts can be attached to the base font, you can effectively use all Unicode characters for which at least one of the fonts contains a suitable glyph.
- ▶ Glyphs from a particular fallback font can be used to override glyphs in the base font, i.e. you can replace glyphs in a font. You can replace one or more individual glyphs, or specify one or more ranges of Unicode characters which will be replaced.
- ▶ The size and vertical position of glyphs from a fallback font can be adjusted to match the base font.

Somewhat surprisingly, the base font itself can also be used as a fallback font. This can be used to implement the following tricks:

- ▶ Adjust the size or position of some or all glyphs in the font.
- ▶ Add characters outside the actual encoding of the base font.

The fallback font mechanism is controlled by the *fallbackfonts* font option, and affects all text output functions. As with all font options the *fallbackfonts* option can be provided in explicit calls to *PDF_load_font()* or in option lists for implicit font loading. Since more than one fallback font can be specified for a base font, the *fallbackfonts* option expects a list of option lists, i.e. an extra set of braces is required.

PDF_info_font() can be used to query the results of the fallback font mechanism (see Section 6.5.3, »Querying Codepage Coverage and Fallback Fonts«, page 133).

Caveats. Note the following when working with fallback fonts:

- ▶ Not all font combinations result in typographically pleasing results. Care should be taken to use only fallback fonts in which the glyph design matches the glyph design of the base font.
- ▶ Font options for the fallback font(s) must be specified separately in the *fallbackfonts* option list.
- ▶ Script-specific shaping (options *shaping*, *script*, *locale*) and OpenType features (options *features*, *script*, *language*) are only applied to glyphs within the same font, but not across glyphs from the base font and a fallback font.
- ▶ The underline/overline/strikeout features must be used with care when working with fallback fonts, as well as *ascender* and similar typographic values. The underline thickness or position defined in the base font may not match the values in the fallback font. As a result, the underline position or thickness may jump in unpleasant ways. A simple workaround against such artifacts is to specify a unified value with the *underlineposition* and *underlinewidth* options of *PDF_fit_textline()* and *PDF_add/create_textflow()*. This value should be selected so that it works with the base font and all fallback fonts.

In the sections below we describe several important use cases of fallback fonts and demonstrate the corresponding option lists.

Add color emoji characters to a text font. Emoji characters are often available in a separate color font (see Section 6.1.3, »OpenType Color Fonts in SVG, COLR and sbix Flavours«, page 110) and require switching between text and emoji fonts. This can be simplified by attaching the emoji font as fallback to the text font. Since the color font may include both mono and color versions of emoji characters we ensure color glyphs with *colormode=ignoremono*:

```
fallbackfonts={
  {fontname={Segoe UI Emoji} colormode=ignoremono}
}
```

Combine fonts for use with multiple scripts. In some situations the script of incoming text data is not known in advance. For example, a database may contain Latin, Greek, and Cyrillic text, but the available fonts cover only one of these scripts at a time. Instead of determining the script and selecting an appropriate font you can construct a font which chains together several fonts, effectively covering the superset of all scripts. Use the following font option list for the *fallbackfonts* option to add Greek and Cyrillic fonts to a Latin font:

```
fallbackfonts={
  {fontname=Times-Greek forcechars={U+0391-U+03F5}}
```

```
{fontname=Times-Cyrillic forcechars={U+0401-U+0490}}  
}
```

Enlarge some or all glyphs in a font. Fallback fonts can also be used to enlarge some or all glyphs in a font without changing the font size. Again, the base font itself will be used as fallback font. This feature can be useful to make different font designs visually compatible without adjusting the font size in the code. Use the following font option list for the *fallbackfonts* option to enlarge all glyphs in the specified range to 120%:

```
fallbackfonts={  
  {fontname=Times-Italic forcechars={U+0020-U+00FF} fontsize=120%}  
}
```

Add an enlarged pictogram. Use the following font option list for the *fallbackfonts* option to pull a symbol from the ZapfDingbats font:

```
fallbackfonts={  
  {fontname=ZapfDingbats forcechars=.a12 fontsize=150% textrise=-15%}  
}
```

Again, we use the *fontsize* and *textrise* suboptions to adjust the symbol's size and position to the base font.

Replace glyphs in a CJK font. You can use the following font option list for the *fallbackfonts* option to replace the Latin characters in the ASCII range with those from another font:

```
fallbackfonts={  
  {fontname=Courier-Bold forcechars={U+0020-U+007E}}  
}
```

Add Latin characters to an Arabic font. This use case is detailed in Section 7.5.4, «Arabic Text Formatting», page 149.

Identify missing glyphs. The font *Unicode BMP Fallback SIL*, which is freely available, displays the hexadecimal value of each Unicode character instead of the actual glyph. This font can be useful for diagnosing font-related problems in the workflow. You can use the following font option list for the *fallbackfonts* option to augment any font with this special fallback font to visualize missing characters:

```
fallbackfonts={{fontname={Unicode BMP Fallback SIL}}}
```

Add Gaiji characters to a font. This use case is detailed in Section 7.6.3, «EUDC and SING Fonts for Gaiji Characters», page 151.

6.4 Font Embedding and Subsetting

6.4.1 Embedding

PDF font embedding and font substitution in Acrobat. PDF documents can include font data in various formats to ensure proper text presentation. Alternatively, a font descriptor containing only the glyph metrics and general font information (but not the

glyph outlines) can be embedded. If a font is not embedded in a PDF document, the PDF viewer takes it from the target system if available and configured («Use local fonts» in Acrobat) or tries to build a substitute font according to the font descriptor. The use of substitution fonts results in readable text, but the glyphs may look different from the original font. Similarly, substitution fonts don't work if complex script shaping or OpenType layout features have been used. For these reasons font embedding is generally recommended unless you know that the documents will be displayed on the target systems acceptably even without embedded fonts. Such PDF files are inherently non-portable, but may be of use in controlled environments, such as corporate networks where the required fonts are known to be available on all workstations. In general font embedding is recommended and is the default behavior in PDFlib.

Embedding fonts with PDFlib. Font embedding is enabled by default, but can be controlled with the *embedding* font option:

```
font = p.load_font("WarnockPro", "winansi", "noembedding");
```

Font embedding for fonts which are exclusively used for invisible text (mainly useful for OCR results) can be controlled with the *optimizeinvisible* font option.

6.4.2 Subsetting

In order to decrease the size of the PDF output, PDFlib can embed only those glyphs of a font which are actually used in the document. This process is called font subsetting. It creates a new font which contains fewer glyphs than the original font, and omits font information which is not required for PDF viewing. Type 3 fonts require special handling for subsetting, see «Type 3 font subsetting», page 131.)

- ▶ When a font for which subsetting has been requested is used in a document, PDFlib keeps track of the glyphs actually used for text output. Subsetting is enabled for fonts which are larger than *subsetminsize* and can be disabled with *subsetting=false*.

Initial font subset. Font subsets contain outline descriptions for all glyphs used in the document. This means that the generated font subsets vary among documents since a different set of glyphs is generally used in each document. Different font subsets can be problematic when merging many small documents with embedded font subsets to a larger document: the embedded subsets cannot be removed since they are all different.

For this scenario PDFlib allows you to specify the initial contents of a font subset with the *initialsubset* option of *PDF_load_font()*. While PDFlib starts with an empty subset by default and adds glyphs as required by the generated text output, the *initialsubset* option can be used to specify a non-empty subset. For example, if you know that only Latin-1 text output will be generated and the font contains many more glyphs, you can specify the first Unicode block as initial subset:

```
initialsubset={U+0020-U+00FF}
```

This means that the glyphs for all Unicode characters in the specified range are included in the subset. If this range is selected so that it covers all text in the generated documents, the generated font subsets will be identical in all documents. As a result, when combining such documents later to a single PDF the identical font subsets can be removed with the *optimize* option of *PDF_begin_document()*.

Type 3 font subsetting. Type 3 fonts must be defined before they can be used in a document (because the glyph widths are required). On the other hand, subsetting is only possible after creating text output (since the glyphs used in the document must be known to create the subset). In order to avoid this conflict, PDFlib supports widths-only Type 3 fonts. If you need subsetting for a Type 3 font you must define the font in two passes; see PDFlib API Reference for details.

Cookbook A full code sample can be found in the *Cookbook* topic `type3_fonts/type3_subsetting`.

6.5 Querying Font Information

`PDF_info_font()` can be used to query useful information related to fonts, encodings, Unicode, and glyphs. Depending on the type of query, a valid font handle may be required as parameter for this method. In all examples below we use the variables described in Table 6.4.

Table 6.4 Variables for use in the examples for `PDF_info_font()`

variable	comments
int uv;	Numerical Unicode value; as an alternative glyph name references without the & and ; decoration can be used in the option list, e.g. unicode=euro. For more details see the description of the Unichar option list data type in the PDFlib API Reference.
int c;	8-bit character code
int gid;	glyph id
int cid;	CID value
String gn;	glyph name
int gn_idx;	String index for a glyph name; if gn_idx is different from -1 the corresponding string can be retrieved as follows: gn = p.get_string(gn_idx, "");
String enc;	encoding name
int font;	valid font handle created with <code>PDF_load_font()</code>

If the requested combination of keyword and option(s) is not available, `PDF_info_font()` returns -1 (in PHP: 0). This must be checked by the client application and can be used to decide whether or not a required glyph is present in a font.

Each of the sample code lines below can be used in isolation since they do not depend on each other.

6.5.1 Font-independent Encoding, Unicode, and Glyph Name Queries

Encoding queries. Encoding queries do not require any valid font handle, i.e. the value -1 (in PHP: 0) can be supplied for the `font` parameter of `PDF_info_font()`. Only glyph names known internally to PDFlib can be supplied in `gn`, but not any font-specific glyph names.

Query the Unicode value of an 8-bit code or a named glyph in an 8-bit encoding:

```
uv = (int) p.info_font(-1, "unicode", "code=" + c + " encoding=" + enc);
uv = (int) p.info_font(-1, "unicode", "glyphname=" + gn + " encoding=" + enc);
```

7Unicode and glyph name queries. `PDF_info_font()` can also be used to perform queries which are independent from a specific 8-bit encoding, but affect the relationship of Unicode values and glyph names known to PDFlib internally. Since these queries are independent from any font, a valid font handle is not required.

Query the Unicode value of an internally known glyph name:

```
uv = (int) p.info_font(-1, "unicode", "glyphname=" + gn + " encoding=unicode");
```

Query the internal glyph name of a Unicode value:

```
gn_idx = (int) p.info_font(-1, "glyphname", "unicode=" + uv + " encoding=unicode");
```

```
/* retrieve the actual glyph name using the string index */  
gn = p.get_string(gn_idx, "");
```

6.5.2 Font-specific Encoding, Unicode, and Glyph Name Queries

The following queries relate to a specific font which must be identified by a valid font handle. The *gn* variable can be used to supply internally known glyph names as well as font-specific glyph names. In all examples below the return value -1 means that the font does not contain the requested glyph.

Query the Unicode value for a code, glyph ID, named glyph, or CID in a font:

```
uv = (int) p.info_font(font, "unicode", "code=" + c);  
uv = (int) p.info_font(font, "unicode", "glyphid=" + gid);  
uv = (int) p.info_font(font, "unicode", "glyphname=" + gn);  
uv = (int) p.info_font(font, "unicode", "cid=" + cid);
```

Query the glyph id for a code, Unicode value, named glyph, or CID in a font:

```
gid = (int) p.info_font(font, "glyphid", "code=" + c);  
gid = (int) p.info_font(font, "glyphid", "unicode=" + uv);  
gid = (int) p.info_font(font, "glyphid", "glyphname=" + gn);  
gid = (int) p.info_font(font, "glyphid", "cid=" + cid);
```

Query the font-specific name of a glyph specified by code, Unicode value, glyph ID, or CID:

```
gn_idx = (int) p.info_font(font, "glyphname", "code=" + c);  
gn_idx = (int) p.info_font(font, "glyphname", "unicode=" + uv);  
gn_idx = (int) p.info_font(font, "glyphname", "glyphid=" + gid);  
gn_idx = (int) p.info_font(font, "glyphname", "cid=" + cid);
```

```
/* retrieve the actual glyph name using the string index */  
gn = p.get_string(gn_idx, "");
```

Check glyph availability. Using `PDF_info_font()` you can check whether a particular font contains the glyphs you need for your application. As an example, the following code checks whether the Euro glyph is contained in a font:

```
/* We could also use "unicode=U+20AC" below */  
if (p.info_font(font, "code", "unicode=euro") == -1)  
{  
    /* no glyph for Euro sign available in the font */  
}
```

Cookbook A full code sample can be found in the *Cookbook* topic `fonts/glyph_availability`.

Alternatively, you can call `PDF_info_textline()` to check the number of unmapped characters for a given text string, i.e. the number of characters in the string for which no appropriate glyph is available in the font. The following code fragment queries results for a string containing a single Euro character (which is expressed with a glyph name reference). If one unmapped character is found this means that the font does not contain any glyph for the Euro sign:

```
String optlist = "font=" + font + " charref";

if (p.info_textline("&euro;", "unmappedchars", optlist) == 1)
{
    /* no glyph for Euro sign available in the font */
}
```

6.5.3 Querying Codepage Coverage and Fallback Fonts

`PDF_info_font()` can also be used to check whether a font is suited for creating text output in a certain language or script, provided the codepage is known which is required for the text. Codepage coverage is encoded in the OS/2 table of the font. Note that it is up to the font designer to decide what exactly it means that a font supports a particular codepage. Even if a font claims to support a specific codepage this does not necessarily mean that it contains glyphs for all characters in this codepage. If more precise coverage information is required you can query the availability of all required characters as demonstrated in Section 6.5.2, »Font-specific Encoding, Unicode, and Glyph Name Queries«, page 132.

Check whether a font supports a codepage. The following fragment checks whether a font supports a particular codepage:

```
String cp="cp1254";

result = (int) p.info_font(font, "codepage", "name=" + cp);

if (result == -1)
    System.err.println("Codepage coverage unknown");
else if (result == 0)
    System.err.println("Codepage not supported by this font");
else
    System.err.println("Codepage supported by this font");
```

Retrieve a list of all supported codepages. The following fragment queries a list of all codepages supported by a TrueType or OpenType font:

```
cp_idx = (int) p.info_font(font, "codepagelist", "");

if (cp_idx == -1)
    System.err.println("Codepage list unknown");
else
{
    System.err.println("Codepage list:");
    System.err.println(p.get_string(cp_idx, ""));
}
```

This creates the following list for the common Arial font:

```
cp1252 cp1250 cp1251 cp1253 cp1254 cp1255 cp1256 cp1257 cp1258 cp874 cp932 cp936 cp949
cp950 cp1361
```

Query fallback glyphs. You can use `PDF_info_font()` to query the results of the fallback font mechanism (see Section 6.3.6, »Fallback Fonts«, page 127, for details on fallback fonts). The following fragment determines the name of the base or fallback font which is used to represent the specified Unicode character:

```
result = p.info_font(basefont, "fallbackfont", "unicode=U+03A3");
/* if result==basefont the base font was used, and no fallback font was required */
if (result == -1)
{
    /* character cannot be displayed with neither base font nor fallback fonts */
}
else
{
    idx = p.info_font(result, "fontname", "api");
    fontname = p.get_string(idx, "");
}
```

7 Text Output

7.1 Text Output Methods

PDFlib supports text output on several levels:

- ▶ Low-level text output with `PDF_show()` and similar functions;
- ▶ Single-line formatted text output with `PDF_fit_textline()`; This method also supports text on a path.
- ▶ Multi-line text formatted output with Textflow (`PDF_fit_textflow()` and related functions); The Textflow formatter can also wrap text inside or outside of vector-based shapes.
- ▶ Text in tables; the table formatter supports Textline and Textflow contents in table cells.

Low-level text output. Functions like `PDF_show()` can be used to place text at a specific location on the page, without using any formatting aids. This is recommended only for applications with very basic text output requirements (e.g. convert plain text files to PDF), or for applications which already have full text placement information (e.g. a driver which converts a page in another format to PDF). The following fragment creates text output with low-level functions:

```
font = p.load_font("Helvetica", "unicode", "");
p.setfont(font, 12);
p.set_text_pos(50, 700);
p.show("Hello world!");
p.continue_text("says Java");
```

Formatted single-line text output with Textlines. `PDF_fit_textline()` creates text output which consists of single lines and offers a variety of formatting features. However, the position of individual Textlines must be determined by the client application. The following fragment creates text output with a Textline. Since font and fontsize can be specified as options, a preceding call to `PDF_load_font()` is not required:

```
p.fit_textline(text, x, y, "fontname=Helvetica fontsize=12");
```

See Section 9.1, »Placing and Fitting Textlines«, page 197, for more information.

Multi-line text output with Textflow. `PDF_fit_textflow()` creates text output with an arbitrary number of lines and can distribute the text across multiple columns or pages. The Textflow formatter supports a wealth of formatting functions. The following fragment creates text output with Textflow:

```
tf = p.add_textflow(tf, text, optlist);
result = p.fit_textflow(tf, llx, lly, urx, ury, optlist);
p.delete_textflow(tf);
```

See Section 9.2, »Multi-Line Textflows«, page 205, for more information.

Text in tables. Textlines and Textflows can also be used to place text in table cells. See Section 9.3, »Table Formatting«, page 225, for more information.

7.2 Font Metrics and Text Styles

7.2.1 Font and Glyph Metrics

Text position. PDFlib maintains the text position independently from the current point for drawing graphics. While the former can be queried via the *textx/texty* options, the latter can be queried via *currentx/currenty*.

Glyph metrics. PDFlib uses the glyph and font metrics system used in PDF which shall be briefly discussed here.

The font size which must be specified by PDFlib users is the minimum distance between adjacent text lines which is required to avoid overlapping character parts. The font size is generally larger than individual characters in a font, since it spans ascender and descender, plus possibly additional space between lines.

The *leading* (line spacing) specifies the vertical distance between the baselines of adjacent lines of text. By default it is set to the value of the font size. The *capheight* is the height of capital letters such as *T* or *H* in most Latin fonts. The *xheight* is the height of lowercase letters such as *x* in most Latin fonts. The *ascender* is the height of lowercase letters such as *f* or *d* in most Latin fonts. The *descender* is the distance from the baseline to the bottom of lowercase letters such as *j* or *p* in most Latin fonts. The descender is usually negative. The values of *xheight*, *capheight*, *ascender*, and *descender* are measured as a fraction of the font size, and must be multiplied with the required font size before being used.

The *gaplen* property is only available in TrueType and OpenType fonts (it will be estimated for other font formats). The *gaplen* value is read from the font file, and specifies the difference between the recommended distance between baselines and the sum of ascender and descender.

PDFlib may have to estimate one or more of these values since they are not guaranteed to be present in the font file. In order to find out whether real or estimated values are used you can call `PDF_info_font()` to query the *xheight* with the option *faked*. The character metrics for a specific font can be queried from PDFlib as follows:

```
font = p.load_font("Times-Roman", "unicode", "");

capheight = p.info_font(font, "capheight", "");
ascender = p.info_font(font, "ascender", "");
descender = p.info_font(font, "descender", "");
xheight = p.info_font(font, "xheight", "");
```

Note The position and size of superscript and subscript cannot be queried with PDFlib.

Cookbook A full code sample can be found in the *Cookbook* topic `fonts/font_metrics_info`.

7.2.2 Kerning

Some character combinations can lead to unpleasant appearance. For example, two characters *V* next to each other can look like a *W*, and the distance between *T* and *e* must be reduced in order to avoid ugly white space. This compensation is referred to as kerning. Many fonts contain comprehensive kerning information which specifies spacing adjustments for critical letter combinations. There are two PDFlib controls for the kerning behavior:

- ▶ By default, kerning information in a font will be read when loading the font. If kerning is not required the *readkerning* option can be set to *false* in *PDF_load_font()*.
- ▶ Kerning for text output must be enabled with the *kerning* text appearance option which is supported by the text output functions.

Temporarily disabling kerning may be useful, for example, for tabular figures when the kerning data contains pairs of figures, since kerned figures wouldn't line up in a table. Note that modern TrueType and OpenType fonts include special figures for this purpose which can be used with the *Tabular Figures* layout feature and the option *features={tnum}*.

Kerning is applied in addition to any character spacing, word spacing, and horizontal scaling which may be active. PDFlib does not impose any limit for the number of kerning pairs in a font.

Kerning is controlled by the font option *readkerning* and the text option *kerning*. By default kerning is enabled.

7.2.3 Text Styles

Simulated bold fonts. PDFlib supports a mechanism for creating artificial bold text for individual text strings via the *fakebold* option. This method simulates a bold font by stroking the glyph outlines; for Type 3 fonts the text will be placed multiply with different offsets. It is strongly recommended to use real bold font designs for emphasis. The *fakebold* option creates text output which is inferior to real bold text, and may inhibit text extraction.

Cookbook A full code sample can be found in the *Cookbook* topic `fonts/simulated_fontstyles`.

Simulated italic fonts. The *italicangle* option can be used to simulate italic fonts when only a regular font is available. This method creates a fake italic font by skewing the regular font by a user-provided angle. Negative values will slant the text clockwise. Be warned that using a real italic or oblique font will result in much more pleasing output. However, if an italic font is not available the *italicangle* option can be used to easily simulate one. This feature may be especially useful for CJK fonts. Typical values for the *italicangle* option are in the range -12 to -15 degrees.

Note PDFlib does not adjust the glyph width to the new bounding box of the slanted glyph. For example, when generated justified text the italicized glyphs may exceed beyond the fitbox.

Tele Vaso

No kerning

Tele Vaso

Kerning applied

Te Va

Character movement caused by kerning

Fig. 7.1 Kerning

Shadow text. PDFlib can create a shadow effect which will generate multiple instances of text where each instance is placed at a slightly different location. Shadow text can be created with the *shadow* option of *PDF_fit_textline()* and *PDF_add/create_textflow()*. The color of the shadow, its position relative to the main text and graphics state options can be specified in suboptions.

Underline, overline, and strikeout text. PDFlib can be instructed to put lines below, above, or in the middle of text. The stroke width of the bar and its distance from the baseline are calculated based on the font's metrics information. In addition, the current values of the horizontal scaling factor and the text matrix are taken into account when calculating the width of the bar. The respective option names for *PDF_set_text_option()* can be used to switch the underline, overline, and strikeout feature on or off, as well as the corresponding options in the text output functions. The *underlineposition* and *underlinewidth* options can be used for fine-tuning.

The *strokecolor* option is used for drawing the bars. The current linecap option is ignored. The *decorationabove* option controls whether or not the line will be drawn on top of or below the text. Aesthetics alert: in most fonts underlining will touch descenders, and overlining will touch diacritical marks atop ascenders.

Cookbook A full code sample can be found in the *Cookbook* topic `text_output/starter_textline`.

Text rendering modes. PDFlib supports several rendering modes which affect the appearance of text. This includes outline text and the ability to use text as a clipping path. Text can also be rendered invisibly which may be useful for placing text on scanned images in order to make the text accessible to searching and indexing, while at the same time assuring it will not be visible directly. The rendering modes are described in the *PDFlib API Reference*, and can be set with the *textrendering* text option.

When stroking text, text state options such as *strokewidth* and *strokecolor* are applied to the glyph outline. The rendering mode has no effect on text displayed using a Type 3 font.

Cookbook Full code samples can be found in the *Cookbook* topics `text_output/text_as_clipping_path` and `text_output/invisible_text`.

Text color. Text is usually displayed in the color specified in the *fillcolor* option. However, if a rendering mode other than `o` has been selected, both *strokecolor* and *fillcolor* affect the text depending on the selected rendering mode.

Cookbook A full code sample can be found in the *Cookbook* topic `text_output/starter_textline`.

7.3 Emoji Sequences

Cookbook *Emoji sequence examples can be found in the Cookbook* topic `text_output/emoji_sequences`.

The appearance of emoji characters can be modified with various sequences, e.g. for changing the skin tone of an emoji character. The underlying concepts are described in Unicode Technical Standard #51 »Unicode Emoji«; comprehensive tables are available at

unicode.org/Public/emoji/14.0/emoji-zwj-sequences.txt
unicode.org/Public/emoji/14.0/emoji-sequences.txt

Since the implementation is based on OpenType feature definitions, support for emoji sequences depends on the font. Modern color fonts such as the freely available *Emoji-OneColor* font or the Windows font *Segoe UI Emoji* support a variety of emoji combinations. Although emoji sequences may also be used in monochrome fonts, the best effects can be achieved with color fonts. Note that the macOS system font *Apple Color Emoji* doesn't include OpenType feature definitions and therefore none of the emoji sequences works with this font.

If a particular sequence is supported by a font, PDFlib's text output methods render the corresponding effect by default without any specific option (the sequences are implemented with the *ccmp* OpenType feature which is enabled by default):

- ▶ Emoji zerowidth joiner sequences for multi-person groupings, role, gender, hair and other modifications using U+200D ZERO WIDTH JOINER (ZWJ). This character can be addressed as character reference `‍` in PDFlib if the text option `charref=true` is set.
- ▶ Skin tone modifiers for adjusting the skin color of emoji using U+1F3FB EMOJI MODIFIER FITZPATRICK TYPE 1-2 - U+1F3FF EMOJI MODIFIER FITZPATRICK TYPE 6. These are usually rendered as color patches if used outside of a skin tone sequence.
- ▶ Emoji flag sequences for creating flags from regional indicator characters U+1F1E6 REGIONAL INDICATOR SYMBOL LETTER A etc.
- ▶ Emoji keycap sequences using digits, U+FE0F EMOJI PRESENTATION SELECTOR (VS16) and U+20E3 COMBINING ENCLOSING KEYCAP

Table 7.1 demonstrates several emoji sequences.

Table 7.1 *Emoji modification sequences*

sequence type	emoji sequence	isolated characters	combined characters
ZWJ grouping	<code>&#x1F468;&zwj;&#x1F469;&zwj;&#x1F467;</code>		
skin tone	<code>&#x1F466;&#x1F3FB;&#x1F466;&#x1F3FD;&#x1F466;&#x1F3FF;</code>		
emoji flag	<code>&#x1F1F2;&#x1F1F9;&#x1F1EC;&#x1F1E7;</code>		
emoji keycap	<code>o&#xFE0F;&#x20E3; 9&#xFE0F;&#x20E3;</code>		
ZWJ gender	<code>&#x1F9DC;&zwj;&#x2642;&#x1F9DC;&zwj;&#x2640;</code>		
skin tone and gender	<code>&#x1F9DC;&#x1F3FF;&zwj;&#x2642;&#x1F9DC;&#x1F3FF;&zwj;&#x2640;</code>		
ZWJ hair components	<code>&#x1F468;&zwj;&#x1F9B0;&#x1F468;&zwj;&#x1F9B3;</code>		
ZWJ role: male	<code>&#x1F468;&zwj;&#x1F33E;</code>		
ZWJ role: female	<code>&#x1F469;&zwj;&#x1F33E;</code>		

7.4 OpenType Layout Features

Cookbook Full code samples can be found in the Cookbook topics `fonts/starter_opentype` and `fonts/opentype_feature_tester`.

7.4.1 Supported OpenType Layout Features

PDFlib supports enhanced text output according to additional information in some fonts. These font extensions are called OpenType layout features. For example, a font may contain a *liga* feature which includes the information that the *f*, *f*, and *i* glyphs can be combined to form a ligature. Other common examples are small caps in the *smcp* feature, i.e. uppercase characters which are smaller than the regular uppercase characters, and old-style figures in the *onum* feature with ascenders and descenders (as opposed to lining figures which are all placed on the baseline). Although ligatures are a very common OpenType feature, they are only one of many features. PDFlib supports the following groups of OpenType features:

- ▶ OpenType features for Western typography are listed in Table 7.2; these are controlled by the *features*, *script* and *language* options.
- ▶ OpenType features for complex script shaping; these are controlled by the *shaping* and *script* options (see Section 7.5, »Complex Script Output«, page 145).
- ▶ OpenType feature tables for kerning. PDFlib doesn't treat *kerning* as OpenType feature because kerning data may also be represented by other means than OpenType feature tables. Use the *readkerning* font option and the *kerning* text option to control kerning (see Section 7.2.2, »Kerning«, page 136).
- ▶ Additional OpenType features for Chinese, Japanese, and Korean text are discussed in Section 7.6.4, »OpenType Layout Features for CJK Text Output«, page 152.

You can use the following Web site to explore the OpenType features in a font:

wakamaifondue.com/

Note Except for kerning PDFlib does not support OpenType features based on the GPOS table. OpenType layout features are not available for 8-bit encodings.

Querying OpenType features. You can query OpenType features in a font programmatically with `PDF_info_font()`. The following statement retrieves a space-separated list with all OpenType features which are available in the font and are supported by PDFlib:

```
result = (int) p.info_font(font, "featurelist", "");
if (result != -1)
{
    /* retrieve string containing space-separated feature list */
    featurelist = p.get_string(result, "");
}
else
{
    /* no supported features found */
}
```

Use the following fragment to check whether PDFlib and the test font support a particular feature, e.g. ligatures (*liga*):

```
result = (int) p.info_font(font, "feature", "name=liga");
if (result == 1)
{
    /* feature supported by font and PDFlib */
}
```

Table 7.2 Supported OpenType features for Western typography and emoji (Table 7.9 lists OpenType features for CJK text)

tag	name	description
ligatures and alternate forms		
aalt	access all alternates	Makes all variations of a selected character accessible. For example, the user may need a glyph outside the context supported by the normal substitution, or the user may not know what feature produces the desired glyph.
calt ¹	contextual alternates	Replace default glyphs with alternate forms which provide better joining behavior. Used in script typefaces which are designed to have some or all of their glyphs join.
clig ¹	contextual ligatures	Replace a sequence of glyphs with a single glyph which is preferred for typographic purposes. Unlike other ligature features, clig specifies the context in which the ligature is recommended. This capability is important in some script designs and for swash ligatures.
cswl	contextual swash	Replace the default glyph with a swash glyph in a specified context.
cv01 ... cv99	character variants 01-99	A font may have stylistic-variant glyphs for one or more characters where the variations for one character are not systematically related to those for other characters. Or, a variation may exist for a character and its casing pair (or related precomposed characters), but not be applicable to other unrelated characters.
dlig	discretionary ligatures	Replace a sequence of glyphs with a single glyph which is preferred for typographic purposes. dlig covers ligatures which may be used for special effects, at the user's preference.
hist	historical forms	Replace the default forms with the historical alternates. Some letter forms were in common use in the past, but appear anachronistic today.
hlig	historical ligatures	Replace the default ligatures with the historical alternates.
liga ¹	standard ligatures	Replace a sequence of glyphs with a single glyph which is preferred for typographic purposes. liga covers ligatures which should be used in normal conditions.
locl ¹	localized forms	Enable localized forms of glyphs to be substituted for default forms. This feature requires the script and language options.
salt	stylistic alternates	Replace the default forms with stylistic alternates. These alternates don't always fit into a clear category like swash or historical.
ss01 ... ss20	stylistic set 01-20	In addition to, or instead of, stylistic alternatives of individual glyphs (see salt feature), some fonts may contain sets of stylistic variant glyphs corresponding to portions of the character set, e.g. multiple variants for lowercase letters in a Latin font.
subs	subscript	Replace a default glyph with a subscript glyph.
sup	superscript	Replace lining or oldstyle figures with superior figures (primarily for footnote indication), and replace lowercase letters with superior letters (primarily for abbreviated French titles)
swsh	swash	Replace default glyphs with swash glyphs.
titl	titling	Replace default glyphs with forms designed for titling.
digits and math		
afrc	alternative fractions	Replace figures separated by a slash with an alternative form.
dnom	denominators	Replace figures which follow a slash with denominator figures.
frac	fractions	Replace figures separated by a slash with common diagonal fractions.
lnum	lining figures	Change figures from oldstyle to the default lining form.
mgrk	mathematical Greek	Replace standard typographic forms of Greek glyphs with forms commonly used in mathematical notation.

Table 7.2 Supported OpenType features for Western typography and emoji (Table 7.9 lists OpenType features for CJK text)

tag	name	description
numr	numerators	Replace figures which precede a slash with numerator figures and replace the typographic slash with the fraction slash.
onum	oldstyle figures	Change figures from the default lining style to oldstyle form.
pnum	proportional figures	Replace monospaced (tabular) figures with figures which have proportional widths.
sinf	scientific inferiors	Replace lining or oldstyle figures with inferior figures (smaller glyphs), primarily for chemical or mathematical notation).
tnum	tabular figures	Replace proportional figures with monospaced (tabular) figures.
zero	slashed zero	Replace the glyph for the figure zero with an alternative form which uses a diagonal slash through the counter.
case and ordinals		
c2pc	petite capitals from capitals	Turn capital characters into petite capitals.
c2sc	small capitals from capitals	Turn capital characters into small capitals.
case	case-sensitive forms	Shift accent marks upwards so that they work better with all-capital sequences or lining figures; also changes oldstyle figures to lining figures.
ordn	ordinals	Replace default alphabetic glyphs with the ordinal forms for use after figures; commonly also creates the Numero character U+2116.
pcap	petite capitals	Turn lowercase characters into petite capitals, i.e. capital letters which are shorter than regular small caps.
smcp	small capitals	Turn lowercase characters into small capitals.
unic	unicase	Map upper- and lowercase letters to a mixed set of lowercase and small capital forms, resulting in a single case alphabet.
positional forms		
falt	final glyph on line alternates	Replace line final glyphs with alternate forms (they would have less or more advance width as need may be), to help justification of text.
isol ² init medi fina	isolated, initial, middle, final forms	Replace glyphs occurring outside a word (isolated), at the beginning, in the middle, or at the end of a word with alternate forms for scripts with joining behavior, especially Arabic.
other		
_none	all features disabled	Deactivate all OpenType features
ccmp ¹	glyph composition/ decomposition	To minimize the number of glyph alternates, it is sometimes desired to decompose a character into two glyphs, or to compose two characters into a single glyph for better glyph processing. This feature is often used for emoji sequences.
ornm	ornaments	Replace the bullet character and/or ASCII characters with ornaments.

1. This feature is enabled by default.

2. These features are enabled automatically if `shaping=true`.

7.4.2 OpenType Layout Features with Textlines and Textflows

PDFlib supports OpenType layout features in the Textline and Textflow functions, but not in the simple text output functions `PDF_show()` etc.

Requirements for OpenType layout features. An OpenType font for use with layout features must contain a GSUB table with the OpenType feature(s) to be used and must be loaded with `encoding=unicode` or `glyphid`.

Note the following when working with OpenType features:

- ▶ OpenType features (options `features`, `script`, `language`) are only applied to glyphs within the same font, but not across glyphs from the master font and one or more fallback fonts if fallback fonts have been specified.
- ▶ Make sure to enable and disable features as you need them. Accidentally leaving OpenType features activated for all text may lead to unexpected results.
- ▶ The following features are enabled by default: `calt`, `ccmp`, `clig`, `liga`, `locl`.

Controlling OpenType features. You can enable and disable OpenType features for fragments of text as required. Use the `features` text option to enable features by name and disable them by prepending `no` to the feature name. For example, use the following to enable the small caps feature with inline option lists for Textflow:

```
<features={smcp}>PDF<features={nosmcp}>
```

For Textlines you can enable OpenType features as follows:

```
p.fit_textline("PDF", x, y, "features={smcp}");
```

You can disable all features with the keyword `_none` or by ignoring the relevant OpenType tables with the font option `readfeatures=false`. OpenType features can also be enabled as Block properties for use with the PDFlib Personalization Server (PPS).

Disabling ligatures with control characters. Some languages disallow the use of ligatures in certain situations. Typographic rules for German and other languages prohibit the use of ligatures across composition boundaries. For example, the `f+i` combination in the word *Schilfinsel* must not be replaced with a ligature since it spans the boundaries between two combined words.

As described above, you can enable and disable ligatures and other OpenType feature processing with the `features` option. Disabling ligatures via options can be cumbersome in cases like the one described above. In order to offer simple ligature control you can disable ligatures with control characters in the text, avoiding the need for enabling/disabling features with multiple options. Inserting the character *Zero-width non-joiner* (U+200C, `&zwnj`; see also Table 7.6) between the constituent characters will prevent them from being replaced by a ligature even if ligatures are enabled. For example, the following sequence will not create any `f+i` ligature:

```
<features={liga charref=true}>Schilf&zwnj;insel
```

7.4.3 Script- and Language-specific OpenType Layout Features

OpenType features may apply in all situations, or can be implemented for a particular script or even a particular script and language combination. For this reason the `script`

and *language* text options can be supplied with the *features* option. They have an effect only if the feature is implemented in a script- or language-specific manner in the font.

As an example, the ligature for the *f* and *i* glyphs is not activated in some fonts if the Turkish language is selected (since the ligated form of *i* could be confused with the dot-less *i* which is used in Turkish). With such a font the following Textflow option list creates an unwanted ligature since no script/language is specified:

```
<features={liga}>fi
```

However, the following Textflow option list does not create any ligature due to the Turkish language option:

```
<script=latn language=TRK features={liga}>fi
```

The *locl* feature selects language-specific character forms, e.g.

Variant character for Serbian:

```
<features={locl} script=cyrl language=SRB charref>&#x0431;
```

Variant figures for Urdu:

```
<features={locl} script=arab language=URD charref>&#x0662;&#x0663;&#x0664;&#x0665;
```

The *script* and *language* options support the keywords defined in the OpenType specification. Examples for various regions of the world are provided in Table 7.3 and Table 7.4. The script keywords are exactly four characters long. Shorter keywords contain trailing space characters which must be included in the value of the *script* option.

Table 7.3 Supported scripts for use with the features option; not all supported scripts are listed.

<i>script name</i>	<i>keyword</i>	<i>script name</i>	<i>keyword</i>
Special		South and Central Asia I: Official Scripts of India	
<i>unspecified script</i>	_none	<i>Devanagari</i>	deva
<i>automatic script detection</i>	_auto	<i>Gujarati</i>	gujr
Europe		<i>Telugu</i>	telu
<i>Latin</i>	latn	<i>Bengali</i>	beng
<i>Cyrillic</i>	cyrl	<i>Odia (formerly Oriya)</i>	orya
<i>Greek</i>	grek	<i>Kannada</i>	knda
Middle East		<i>Gurmukhi</i>	guru
<i>Hebrew</i>	hebr	<i>Tamil</i>	taml
<i>Arabic</i>	arab	<i>Malayalam</i>	mlym
East Asia		South and Central Asia II: Other Modern Scripts	
<i>CIK Ideographic</i>	hani	<i>Tibetan</i>	tibt
<i>Hangul</i>	hang	<i>Mongolian</i>	mong
<i>Hangul Jamo</i>	jamo	Southeast Asia	

Table 7.3 Supported scripts for use with the features option; not all supported scripts are listed.

script name	keyword	script name	keyword
Bopomofo	bopo	Thai	thai
Yi	yi	Lao	lao
Tangut	tang	Tai Viet	tavt
Hiragana	kana	Notational systems	
Katakana	kana	Braille	brai

Table 7.4 Keywords for the language option; not all supported languages are listed.

key-word	language	key-word	language	key-word	language
_none	unspecified language	ELL	Greek	PLK	Polish
ARA	Arabic	GUJ	Gujarati	ROM	Romanian
BEN	Bengali	IWR	Hebrew	RUS	Russian
BGR	Bulgarian	HIN	Hindi	SRB	Serbian
CAT	Catalan	HUN	Hungarian	SND	Sindhi
ZHP	Chinese phonetic	IND	Indonesian	SNH	Sinhalese
ZHS	Chinese simplified	JAN	Japanese	ESP	Spanish
ZHT	Chinese traditional	KAN	Kannada	TAM	Tamil
NLD	Dutch	KSH	Kashmiri	TEL	Telugu
ENG	English	KHM	Khmer	THA	Thai
FAR	Farsi	KOR	Korean	TIB	Tibetan
FRA	French	ORI	Oriya	TRK	Turkish ¹
DEU	German	PAS	Pashto	URD	Urdu

1. Some fonts wrongly use TUR for Turkish; PDFlib treats this tag as equivalent to TRK.

7.5 Complex Script Output

Cookbook A full code sample can be found in the *Cookbook* topic `complex_scripts/starter_shaping`.

7.5.1 Complex Scripts

The Latin script places one character after the other in left-to-right order. Other scripts (writing systems) have additional requirements for text output. We refer to such writing systems as complex scripts. PDFlib performs text processing for a variety of current and ancient scripts. In this section we discuss this process in more detail. Some scripts require additional processing:

- ▶ Arabic and Hebrew place text from right to left. Mixed text (e.g. Arabic with a Latin insert) contains both right-to-left and left-to-right segments. These segments must be reordered, which is referred to as the Bidi (bidirectional) problem.
- ▶ Some scripts, especially Arabic, use different character shapes depending on the position of the character (isolated, beginning/middle/end of a word).

- ▶ Mandatory ligatures replace sequences of characters.
- ▶ The position of glyphs must be adjusted horizontally and vertically.
- ▶ Indic scripts require reordering of some characters, i.e. characters may change their position in the text.

Shaping. Scripts which require one or more of these processing steps are called complex scripts. The process of preparing incoming logical text for proper presentation is called shaping. Shaping selects appropriate glyph forms depending on whether a character is located at the start, middle, or end of a word, or in a standalone position. Shaping is a crucial component of Arabic and Hindi text formatting. Shaping may also replace a sequence of two or more characters with a suitable ligature or reorder the characters. Since the shaping process determines the character forms automatically, explicit ligatures and Unicode presentation forms (e.g. Arabic Presentation Forms A starting at U+FB50) must not be used as input characters.

Caveats. Note the following when working with complex script shaping:

- ▶ PDFlib does not automatically set the *shaping* and *script* options, but expects them to be supplied by the user.
- ▶ Script-specific shaping (options *shaping*, *script*, *language*) is only applied to glyphs within the same font, but not across glyphs from different fonts. If fallback fonts are used, shaping is only applied within text runs in the same (master or fallback) font.
- ▶ Since shaping may reorder characters in the text, care must be taken regarding attribute changes within a word. For example, if you use inline options in Textflow to colorize the second character in a word – what should happen when shaping swaps the first and second characters? For this reason, formatting changes in shaped text should only be applied at word boundaries, but not within words.

Requirements for shaping. A font for use with complex script shaping must meet the following requirements:

- ▶ The font must contain the necessary OpenType feature tables for the selected script.
- ▶ The font must be loaded with *encoding=unicode* or *glyhid*.
- ▶ The *vertical* option of *PDF_load_font()* must not be used, and the *readshaping* option must not be set to *false*.

7.5.2 Controlling the Shaping Process

Script selection. Complex script processing (option *shaping*) requires the *script* option. The *language* option can additionally be supplied. It controls language-specific aspects of shaping, e.g. different figures for Arabic vs. Urdu. However, only few fonts contain language-specific shaping tables, so in most cases specifying the *script* option will be sufficient, and shaping cannot be improved with the *language* option. The shaping process supports the scripts listed in Table 7.5. The following option list enables shaping for Arabic text (including Bidi processing):

```
shaping script=arab
```

Table 7.5 Supported scripts for use with the shaping option

region	supported scripts and corresponding keywords
Middle East	Arabic (arab), Hebrew (hebr)

Table 7.5 Supported scripts for use with the shaping option

region	supported scripts and corresponding keywords
Central Asia	Sinhala (sinh), Tibetan (tibt)
Southeast Asia	Khmer (khmr), Thai (thai)
East Asia	Hangul (hang), CJK Ideographic (hani)
India	Bengali (beng), Devanagari (deva), Gujarati (gujr), Kannada (knda), Malayalam (mlym), Odia, formerly Oriya (orya), Gurmukhi (guru), Tamil (tam1), Telugu (telu)

With `script=_auto` PDFlib automatically assigns the script to which the majority of characters in the text belong. Since Latin text doesn't require shaping it is ignored when determining the script automatically. You can query the scripts used for some text with the `scriptlist` keyword of `PDF_info_textline()`.

Language selection. The `language` option specifies the natural language in which the text is written. It supports the keywords listed in Table 7.4. Examples:

```
language=ARA
language=URD
```

Note While advanced line breaking for Textflow (see Section 9.2.5, »Advanced script-specific Line Breaking«, page 221) also applies language-specific processing, it is not controlled by the `language` option, but by the `locale` option which identifies not only languages, but also countries and regions.

Overriding shaping behavior. In some cases you may want to override the shaping behavior. PDFlib supports several Unicode formatting characters for this purpose. For convenience, these formatting characters can also be specified with entity names (see Table 7.6).

Table 7.6 Unicode control characters for overriding the default shaping behavior

formatting character	entity name	Unicode name	function
U+200C	ZWNJ	ZERO WIDTH NON-JOINER	prevent the two adjacent characters from connecting
U+200D	ZWJ	ZERO WIDTH JOINER	force the two adjacent characters to connect

7.5.3 Bidirectional Formatting

Cookbook A full code sample can be found in the Cookbook topic `complex_scripts/bidi_formatting`.

For right-to-left text (especially Arabic and Hebrew, but also some other scripts) it is very common to have nested sequences of left-to-right Latin text, e.g. an address or a quote in another language. These mixed sequences of text require bidirectional (Bidi) formatting. Since numerals are always written from left to right, the Bidi problem affects even text which is completely written in Arabic or Hebrew. Bidi processing does not have to be enabled with an option, but will automatically be applied as part of the shaping process if text in a right-to-left script with an appropriate `script` option is encountered.

Note Bidi processing is not supported for multi-line Textflows, but only for Textlines (i.e. single-line text output).

Overriding the Bidi algorithm. While automatic Bidi processing will provide proper results in common cases, there are situations which require explicit user control. PDFlib supports several directional formatting codes for this purpose. For convenience, these formatting characters can also be specified with entity names (see Table 7.7). The bidirectional formatting codes are useful to override the default Bidi algorithm in the following situations:

- ▶ a right-to-left paragraph begins with left-to-right characters;
- ▶ there are nested segments with mixed text;
- ▶ there are weak characters, e.g. punctuation, at the boundary between left-to-right and right-to-left text;
- ▶ part numbers and similar entities containing mixed text.

Table 7.7 Directional formatting codes for overriding the bidirectional algorithm

formatting code	entity name	Unicode name	function
U+202A	LRE	LEFT-TO-RIGHT EMBEDDING (LRE)	start an embedded left-to-right sequence
U+202B	RLE	RIGHT-TO-LEFT EMBEDDING (RLE)	start an embedded right-to-left sequence
U+200E	LRM	LEFT-TO-RIGHT MARK (LRM)	left-to-right zero-width character
U+200F	RLM	RIGHT-TO-LEFT MARK (RLM)	right-to-left zero-width character
U+202D	LRO	LEFT-TO-RIGHT OVERRIDE (LRO)	force characters to be treated as strong left-to-right characters
U+202E	RLO	RIGHT-TO-LEFT OVERRIDE (RLO)	force characters to be treated as strong right-to-left characters
U+202C	PDF	POP DIRECTIONAL FORMATTING (PDF)	restore the bidirectional state to what it was before the last LRE, RLE, RLO, or LRO

Options for improved right-to-left document handling. The default settings of various formatting options and Acrobat behavior are targeted at left-to-right text output. Use the following options for right-to-left text formatting and document display:

- ▶ Place a Textline right-aligned with the following fitting option:
`position={right center}`
- ▶ Create a leader between the text and the left border:
`leader={alignment=left text=.`
- ▶ Use the following option of `PDF_begin/end_document()` to activate better right-to-left document and page display in Acrobat:
`viewerpreferences={direction=r2l}`

Dealing with Bidi text in your code. The following may also be useful when dealing with bidirectional text:

- ▶ You can use the `startx/starty` and `endx/endy` keywords of `PDF_info_textline()` to determine the coordinates of the logical start and end characters, respectively.

- ▶ You can use the *writingdirx* keyword of `PDF_info_textline()` to determine the dominant writing direction of text. This direction will be inferred from the initial characters of the text or from directional formatting codes according to Table 7.7 (if present in the text).
- ▶ You can use the *auto* keyword for the *position* option of `PDF_info_textline()` to automatically align Arabic or Hebrew text at the right border and Latin text at the left border. For example, the following Textline option list aligns right- or left-aligns the text on the baseline:

```
boxsize={width 0} position={auto bottom}
```

7.5.4 Arabic Text Formatting

Cookbook A full code sample can be found in the Cookbook topic `complex_scripts/arabic_formatting`.

In addition to Bidirectional formatting and text shaping as discussed above there are several other formatting aspects related to generating text output in the Arabic script.

Arabic ligatures. The Arabic script makes extensive use of ligatures. Many Arabic fonts contain two kinds of ligatures which are treated differently in PDFlib:

- ▶ Required ligatures (*rlig* feature) must always be applied, e.g. the Lam-Alef ligature and variants thereof. Required ligatures are automatically used if the *shaping* option is enabled with `script=arab`.
- ▶ Optional Arabic ligatures (*liga* and *dlig* features) are not used automatically, but can be enabled like other user-controlled OpenType features, i.e. `features={dlig}`. Optional Arabic ligatures are applied after complex script processing and shaping.

Avoiding ligatures. In some cases joining adjacent characters is not desired, e.g. for certain abbreviations. In such cases you can use the formatting characters listed in Table 7.6 to force or prevent characters from being joined. For example, the zero-width non-joiner character in the following example prevents the characters from being joined in order to form a proper abbreviation:

```
&#x0623;&#x064A;&ZWJ;&#x0628;&#x064A;&ZWJ;&#x0625;&#x0645;
```

Tatweel formatting for Arabic text. You can stretch Arabic words by inserting one or more instances of the tatweel character U+0640 (also called kashida). While PDFlib does not automatically justify text by inserting tatweel characters, you can insert this character in the input text to stretch words.

Adding Latin characters to an Arabic font. Some Arabic fonts do not contain any glyphs for Latin characters, e.g. Google's Noto fonts. In this situation you can use the *fallbackfonts* option to merge Latin characters into an Arabic font. PDFlib automatically switches between both fonts depending on the Latin or Arabic text input, i.e. you don't have to switch fonts in your application but can supply mixed Latin/Arabic text with a single font specification.

You can use the following font loading option list for the *fallbackfonts* option to add Latin characters from the NotoSerif-Regular font to the NotoNaskhArabic-Regular font:

```
fontname=NotoNaskhArabic-Regular fallbackfonts={ {fontname=NotoSerif-Regular} }
```

7.6 Chinese, Japanese, and Korean Text Output

7.6.1 Using TrueType and OpenType CJK Fonts

PDFlib supports CJK fonts in the TrueType, TrueType/OpenType Collection (TTC, OTC) and OpenType formats. CJK fonts are processed as follows:

- ▶ If the *embedding* option is *true*, the font is converted to a CID font and embedded in the PDF output.
- ▶ CJK host font names on Windows can be supplied to *PDF_load_font()* as UTF-8 with BOM or UTF-16. Non-Latin host font names are not supported on macOS.

The following example uses the *ArialUnicodeMS* font to display Chinese text. The font must either be installed on the system or must be configured according to Section 6.3.4, »Searching for Fonts«, page 122):

```
font = p.load_font("Arial Unicode MS", "unicode", "");
if (font == -1) { ... }

p.fit_textline("\u4e00\u500b\u4eba", x, y, "fontsize=24");
```

Accessing individual fonts in a TrueType Collection. TTC/OTC files contain multiple separate fonts. You can access each font by supplying its proper name. However, if you don't know which fonts are contained in a TTC/OTC file you can numerically address each font by appending a colon character and the number of the font within the TTC/OTC file (starting with 0). If the index is 0 it can be omitted. For example, the TTC file *msgothic.ttc* contains multiple fonts which can be addressed as follows in *PDF_load_font()* (each line contains equivalent font names):



```
msgothic:0      MS Gothic      msgothic:
msgothic:1      MS PGothic
msgothic:2      MS UI Gothic
```

Note that *msgothic* (without any suffix) does not work as a font name since it does not uniquely identify a font. Font name aliases (see »Sources of Font Data«, page 122) can be used in combination with TTC/OTC indexing. If a font with the specified index cannot be found, the method call fails.

It is only required to configure the TTC/OTC font file once; all indexed fonts in the TTC/OTC file will be found automatically. The following code is sufficient to configure all indexed fonts in *msgothic.ttc* (see Section 6.3.4, »Searching for Fonts«, page 122):

```
p.set_option("FontOutline={msgothic=msgothic.ttc}");
```

7.6.2 Horizontal and Vertical Writing Mode.

PDFlib supports horizontal and vertical writing modes. Vertical writing mode can be requested in different ways:

- ▶ TrueType and OpenType Fonts can be used for vertical writing mode by supplying the *vertical* font option.
- ▶ Font names starting with an '@' character are always processed in vertical mode.

By default all glyphs have the same height in vertical writing mode. However, TrueType and OpenType fonts may contain proportional metrics for vertical writing mode.

PDFlib can be instructed to use such proportional vertical metrics by setting the font option *readverticalmetrics* (with the default *false*).

Note Character spacing must be negative to spread characters apart in vertical writing mode.

OpenType fonts may contain the OpenType layout features listed in Table 7.8. They replace default glyph forms with variants adjusted for vertical writing. The *vkna* feature can be controlled with the *features* text option, while *vrt2/vert* are automatically enabled in vertical writing mode.

Table 7.8 OpenType layout features for vertical text

key-word	name	description
vert ¹	vertical writing	Replace default forms with variants adjusted for vertical writing.
vkna	vertical Kana alternates	Replace standard Kana with forms that have been designed for vertical writing.
vrt2 ¹	vertical alternates and rotation	Replace some fixed-width (half-, third- or quarter-width) or proportional-width glyphs (mostly Latin or Katakana) with forms suitable for vertical writing, i.e. rotated 90° clockwise. If this feature is present it disables the <i>vert</i> feature which is a subset of <i>vrt2</i> .

1. Automatically enabled for fonts in vertical writing mode

7.6.3 EUDC and SING Fonts for Gaiji Characters

PDFlib supports Windows EUDC (end-user defined characters, **.tte*) and SING fonts (**.gai*) which can be used to access custom Gaiji characters for CJK text. Most conveniently fonts with custom characters are integrated into other fonts with the fallback font mechanism.

Using fallback fonts for Gaiji characters. Typically, Gaiji characters are pulled from Windows EUDC fonts or SING glyphlets, but the *fallbackfonts* option accepts any kind of font. Therefore this approach is not limited to Gaiji characters, but can be used for any kind of symbol (e.g. a company logo in a separate font). You can use the following font option list for the *fallbackfonts* option to add a user-defined (gaiji) character from an EUDC font to the loaded font:

```
fallbackfonts={
  {fontname=EUDC forcechars=U+E000 fontsize=140% textrise=-20%}
}
```

Once a base font has been loaded with this fallback font configuration, the EUDC character can be used within the text without any need to change the font.

With SING fonts the Unicode value doesn't have to be supplied since it is automatically determined by PDFlib:

```
fallbackfonts={
  {fontname=PDFlibWing forcechars=gaiji}
}
```

Preserving PUA values for Gaiji characters. In some cases, e.g. with Windows EUDC fonts, Gaiji characters are mapped to Unicode values in the Private Use Area (PUA) by the font. By default, PDFlib replaces PUA values with U+FFFD (Unicode replacement

character) in the generated PDF (ToUnicode CMap). As a result such characters cannot be extracted correctly from the generated PDF.

This behavior can be changed with the *preserve_pua* font option. If it is set to *true*, Gaiji characters with PUA values retain their Unicode values.

7.6.4 OpenType Layout Features for CJK Text Output

As detailed in Section 7.4, »OpenType Layout Features«, page 139, PDFlib supports advanced typographic layout tables in OpenType and TrueType fonts. For example, OpenType features can be used to select alternative forms of the Latin glyphs with proportional widths or half widths, or to select alternate character forms. Table 7.9 lists OpenType features for CJK text (additional OpenType layout features for general use are listed in Table 7.2 and features for vertical writing mode in Table 7.8).

Table 7.9 Supported OpenType layout features for Chinese, Japanese, and Korean text

key-word	name	description
expt	expert forms	Like the JIS78 forms this feature replaces standard Japanese forms with corresponding forms preferred by typographers.
fwid	full widths	Replace glyphs set on other widths with glyphs set on full (usually em) widths. This may include Latin characters and various symbols.
hkna	horizontal Kana alternates	Replace standard Kana with forms that have been designed for only horizontal writing.
hngl	Hangul	(Deprecated per ISO 14496-22:2015/Amd.2:2017) Replace hanja (Chinese-style) Korean characters with the corresponding Hangul (syllabic) characters.
hojo	Hojo Kanji forms (JIS X 0212-1990)	Access the JIS X 0212-1990 glyphs (also called »Hojo Kanji«) if the JIS X 0213:2004 form is encoded as default.
hwid	half widths	Replace glyphs on proportional widths, or fixed widths other than half an em, with glyphs on half-em (widths).
ital	italics	Replace the Roman glyphs with the corresponding Italic glyphs.
jp04	JIS2004 forms	(Subset of the n1ck feature) Access the JIS X 0213:2004 glyphs.
jp78	JIS78 forms	Replace default (JIS90) Japanese glyphs with the corresponding forms from JIS C 6226-1978 (JIS78).
jp83	JIS83 forms	Replace default (JIS90) Japanese glyphs with the corresponding forms from JIS X 0208-1983 (JIS83).
jp90	JIS90 forms	Replace Japanese glyphs from JIS78 or JIS83 with the corresponding forms from JIS X 0208-1990 (JIS90).
loc1 ¹	localized forms	Enable localized forms of glyphs to be substituted for default forms. This feature requires the script and language options.
nalt	alternate annotation forms	Replace default glyphs with various notational forms (e.g. glyphs placed in open or solid circles, squares, parentheses, diamonds or rounded boxes).
n1ck	NLC Kanji forms	Access the new glyph shapes defined in 2000 by the National Language Council (NLC) of Japan for a number of JIS characters.
pkna	proportional Kana	Replace glyphs, Kana and Kana-related, set on uniform widths (half or full-width) with proportional glyphs.

Table 7.9 Supported OpenType layout features for Chinese, Japanese, and Korean text

key-word	name	description
pwid	proportional widths	Replace glyphs set on uniform widths (typically full or half-em) with proportionally spaced glyphs.
qwid	quarter widths	Replace glyphs on other widths with glyphs set on widths of one quarter of an em.
ruby	Ruby notation forms	Replace default Kana glyphs with smaller glyphs designed for use as (usually superscripted) Ruby.
smp1	simplified forms	Replace traditional Chinese or Japanese forms with the corresponding simplified forms.
tnam	traditional name forms	Replace simplified Japanese Kanji forms with the corresponding traditional forms. This is equivalent to the trad feature, but limited to the traditional forms considered proper for use in personal names.
trad	traditional forms	Replace simplified Chinese Hanzi or Japanese Kanji forms with the corresponding traditional forms.
twid	third widths	Replace glyphs on other widths with glyphs set on widths of one third of an em.

1. This feature is enabled by default.

7.7 Unicode Variation Sequences (UVS)

A particular Unicode character can be represented by a variety of different glyphs. Generally, such visual differences are realized by using suitable fonts (e.g. regular vs. italic) or alternative glyphs within a font (e.g. swash characters) which are selected by OpenType features. In some situations the choice of glyph is semantically relevant and must be explicit even in plain text without any font-related formatting information. Unicode offers Variation Selectors for this purpose.

7.7.1 Types of Variation Sequences

A Unicode Variation Sequence (UVS) consists of a base Unicode character followed by a variation selector. The sequence selects a *variant* of the base character. The Unicode standard specifies different types of sequences. OpenType fonts may include a UVS table which specifies variant glyphs and the corresponding variation selectors. The UVS table specifies which glyph variant is used by default and which alternative glyphs are available in the font.

Standardized variation sequences. These use one of the variation selectors *U+FE00 VS1* - *U+FE0D VS14* (currently only *VS1* is standardized) which must be placed after the base character. Standardized variation sequences are used for alternate mathematical glyphs and Mongolian text.

Emoji variation sequences. Several hundred emoji characters are included in a list called *emoji variation sequences*. They are available in text style (usually monochrome) and emoji style (usually in color). This choice can be overridden by placing one of two variation selectors after the Unicode character:

- ▶ *U+FE0E VS15 text presentation selector* selects the text style (monochrome) glyph. For example, the text style may be preferable if you want to apply custom color to the glyph instead of relying on its intrinsic color.

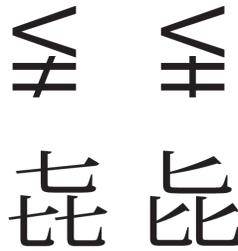


Fig. 7.2
Default glyph (left) and variant glyph (right)

- ▶ *U+FE0F VS16 emoji presentation selector* selects the emoji style (color) glyph.

The Unicode standard specifies the default style of each emoji character, but the UVS table in a font may change the default for individual characters. More importantly, the availability of these variations is up to the font designer.

Ideographic Variation Sequences (IVS). These are defined by a registration process according to Unicode Technical Standard #37 and are listed in the Ideographic Variation Database. For example, variant glyphs are available for CJK compatibility ideographs. An IVS consists of a unified ideographic character as base character and one of 240 variation selectors *U+E0100 VS17 - U+E01EF VS256*. IVSes are mainly used for selecting appropriate glyphs for person and place names.

7.7.2 Variation Sequences in PDFlib

Selecting variant glyphs. PDFlib interprets the UVS table in an OpenType font unless this has been disabled with the font option *readselectors=false*.

Assuming you know that a font contains the required glyphs, working with variation sequences is as simple as providing the sequence to PDFlib's text output functions. The code fragments below print the default glyph of a Unicode base character plus a variant chosen by a selector. The resulting output is shown in figure 7.2; note the differences within the glyph pairs:

```
p.fit_textline("&#x2268; &#x2268;&#xFE00;", 50, 700,  
             "fontname={Cambria Math} fontsize=24 charref=true");  
  
p.fit_textline("&#x3402;&#xE0100; &#x3402;&#xE0101;", 50, 650,  
             "fontname={KozMinPr6N-Regular} fontsize=24 charref=true");
```

If the variation selector for the base character cannot be honored, e.g. because the font doesn't contain the required glyph, it is ignored.

Querying variation selectors in a font. Using *PDF_info_font()* you can check whether a font contains variation selectors at all. The *selector* keyword along with the *index* option retrieves a list of all variation selectors available in the font:

```
for (i = 0; i < 256; i++)  
{  
    selectors[i] = (int) p.info_font(font, "selector", "index=" + i);  
    if (selectors[i] == -1)
```

```

    {
        selectorcount = i;
        break;
    }
}

```

The following code fragment can be used to check whether a font contains a variant glyph for a particular sequence with selector *s* following Unicode value *uv*:

```

if (p.info_font(font, "unicode", "unicode=" + uv + " selector=" + s) != -1)
{
    /* variant glyph for this sequence is available in the font */
}

```

This query is only intended to check whether a variant is available. The resulting Unicode (if a variant is available) itself is unlikely to be useful because PDFlib assigns PUA Unicode values to variants.

7.8 The Text Processing Pipeline

The client application provides text for page output to PDFlib. This text is encoded according to some encoding and format. However, PDFlib's internal processing is based on the Unicode standard, and the final text output requires font-specific glyph IDs. PDFlib therefore treats incoming strings for page contents in a text processing pipeline with three sections:

- ▶ normalize input codes to Unicode values subject to the selected encoding.
- ▶ convert Unicode values to font-specific glyph IDs subject to the available glyphs in the font.
- ▶ transform glyph IDs subject to the output encoding.

Several aspects of the text processing pipeline can be controlled by options.

7.8.1 Normalizing Input Strings to Unicode

The following steps are performed for all encodings except *glyphid*:

- ▶ Unicode-capable language bindings: if a single-byte encoding has been specified UTF-16 text is converted to single-byte text by dropping the high-order byte.
- ▶ Replace escape sequences (see Section 5.4.1, »Escape Sequences«, page 105) with the corresponding numerical values.
- ▶ Resolve character references and replace them with the corresponding Unicode values (see Section 5.4.2, »Character References«, page 106, and next section below).
- ▶ Single-byte encodings: convert single-byte text to Unicode according to the specified encoding.
- ▶ Normalize the text to one of the Unicode normalization forms (e.g. NFC) according to the *normalize* option.

Character references with glyph names. If a glyph name reference (see Section 5.4.2, »Character References«, page 106) is used, PDFlib tries to locate the glyph with the specified name in the font and replaces the reference with the glyph's Unicode value. If a glyph with the specified name is not available in the font, PDFlib searches its internal glyph name table to determine a Unicode value. This Unicode value is used to check

whether a suitable glyph is present in the font. If no such glyph can be found the behavior is controlled by the *glyphcheck* and *errorpolicy* settings. Character references cannot be used with *glyphid* or *builtin* encoding.

Unicode control characters. Control characters are Unicode values which do not represent any glyph, but are used to convey formatting information. PDFlib processes the following groups of Unicode control characters:

- ▶ The control characters for overriding the default shaping behavior (listed in Table 7.6) and those for overriding the default bidi formatting (listed in Table 7.7) control complex script shaping and OpenType layout feature processing in Textline and Textflow. After evaluating these control characters they will be removed.
- ▶ The formatting control characters for line breaking and Textflow formatting listed in Table 9.1. After evaluating these control characters they will be removed.
- ▶ Other Unicode control characters in the ranges U+0001-U+0019 and U+007F-U+009F are replaced with the *replacementchar* character.

Even if a font contains a glyph for a control character the glyph will usually not be visible since PDFlib removes control characters (as an exception to this rule *&NBSP;* and *&SHY;* are not removed). However, with *encoding=glyphid* control characters are retained in the text and can produce visible output.

7.8.2 Converting Unicode Values to Glyph IDs

The Unicode values determined in the previous section may have to be modified for several reasons. The steps below are performed for all encodings with the following exception: with *encoding=glyphid:invalid* glyph IDs are replaced with glyph ID 0. If *glyphcheck=error* an exception is thrown.

Force characters from fallback fonts. Replace Unicode values according to the *forcechars* suboption of the *fallbackfonts* option, and determine the glyph ID of the corresponding fallback font. For more information see Section 6.3.6, »Fallback Fonts«, page 127.

Resolve variation sequences. For some fonts Unicode characters may be followed by a variation selector which selects a specific glyph variant of the character (see Section 7.7, »Unicode Variation Sequences (UVS)«, page 153). If the font contains a variant glyph for the variation sequence the glyph ID of the variant glyph is used instead of the original glyph ID.

Convert to glyph IDs. Convert the Unicode values to glyph IDs according to the mappings determined in Section 6.2.2, »Unicode Mapping for Glyphs«, page 114. If no corresponding glyph ID for a Unicode value was found in the font, the next steps depend on the *glyphcheck* option:

- ▶ *glyphcheck=none*: glyph ID 0 is used, i.e. the *.notdef* glyph is used in the text output. If the *.notdef* glyph contains a visible shape (often a hollow or crossed-out rectangle) it makes the problematic characters visible on the PDF page, which may or may not be desired.
- ▶ *glyphcheck=replace* (which is the default): a warning message is logged and PDFlib attempts to replace the unmappable Unicode value with the glyph replacement mechanism detailed below.

- ▶ *glyphcheck=error*: PDFlib raises an error. In case of *errorpolicy=return* this means that the method call terminates without creating any text output; *PDF_add/create_textflow()* return -1 (in PHP: 0). In case of *errorpolicy=exception* an exception is thrown.

Glyph replacement. If *glyphcheck=replace*, unmappable Unicode values are recursively replaced as follows:

- ▶ The fallback fonts specified when loading the master font are searched for glyphs for the Unicode value. This may involve an arbitrary number of fonts since more than one fallback font can be specified for each font. If a glyph is found in one of the fallback fonts it is used.
- ▶ Select a typographically similar glyph according to the Unicode value from PDFlib's internal replacement table. The following excerpt from the internal list contains some of these replacements. If the first character in the list is unavailable in a font, it is replaced with the second character:

```
U+00A0 (NO-BREAK SPACE)U+0020 (SPACE)
U+00AD (SOFT HYPHEN)U+002D (HYPHEN-MINUS)
U+2010 (HYPHEN)U+002D (HYPHEN-MINUS)
U+03BC (GREEK SMALL LETTER MU)U+00C5 (MICRO SIGN)
U+212B (ANGSTROM SIGN)U+00B5 (LATIN CAPITAL LETTER A WITH RING ABOVE Å)
U+220F (N-ARY PRODUCT)U+03A0 (GREEK CAPITAL LETTER PI)
U+2126 (OHM SIGN)U+03A9 (GREEK CAPITAL LETTER OMEGA)
```

In addition to the internal table, the fullwidth characters U+FF01 to U+FF5E are replaced with the corresponding ISO 8859-1 characters (i.e. U+0021 to U+007E) if the fullwidth variants are not available in the font.

- ▶ Decompose Unicode ligatures into their constituent glyphs (e.g. replace U+FB00 *Latin small ligature ff* with the sequence U+0066 *f*, U+0066 *f*).
- ▶ Select glyphs with the same Unicode semantics according to their glyph name. In particular, all glyph name suffixes separated with a period are removed if the corresponding glyph is not available (e.g. replace *A.swash* with *A*; replace *g.alt* with *g*).

If none of these methods delivers a glyph for the Unicode value, the *replacementchar* option is evaluated as follows:

- ▶ If *replacementchar=auto* (which is the default) the characters U+00A0 (NO-BREAK SPACE) and U+0020 (SPACE) are tried. If these are still unavailable, the »missing glyph« symbol is used (not allowed in PDF/A, PDF/UA and PDF/X-4/5).
- ▶ If a Unicode character was specified as *replacementchar* it is used instead of the original character.
- ▶ If *replacementchar=drop*, the character is dropped from the input stream and no output is created.
- ▶ If *replacementchar=error* an exception is thrown. This may be used to avoid unreadable text output.

Cookbook A full code sample can be found in the *Cookbook topic* `fonts/glyph_replacement`.

7.8.3 Transforming Glyph IDs

The determined glyph IDs are not yet final since several transformations may have to be applied before final output can be created. The details of these transformations depend on the font and several options.

Vertical glyphs. For fonts in vertical writing mode some glyphs may be replaced by their vertical counterparts. This substitution requires a *vert* or *vrt2* OpenType layout feature table in the font.

OpenType layout features. OpenType features can create ligatures, swash characters, small caps, and many other typographic variations by replacing one or more glyph IDs with other values. This topic is discussed in Section 7.4, »OpenType Layout Features«, page 139. OpenType layout features are relevant only for suitable fonts (see »Requirements for OpenType layout features«, page 143), and are applied according to the *features* option. Some layout features are enabled by default.

Complex script shaping. Shaping reorders the text and determines the appropriate variant glyph according to the position of a character (e.g. initial, middle, final, or isolated form of Arabic characters). Shaping is discussed in Section 7.5, »Complex Script Output«, page 145. It is relevant only for suitable fonts (see »Requirements for shaping«, page 146), and is applied according to the *shaping* option.

8 Importing Images, SVG Graphics and PDF Pages

8.1 Raster Images

8.1.1 Basic Image Handling

Embedding raster images with PDFlib is easy to accomplish. First, the image file has to be opened with a PDFlib method which analyzes the image characteristics and copies the pixel data to the PDF output. `PDF_load_image()` returns a handle which serves as an image descriptor. This handle can be used in a call to `PDF_fit_image()`, along with positioning and scaling options:

```
image = p.load_image("auto", "image.jpg", "");
if (image == -1)
    throw new Exception("Error: " + p.get_errmsg());

p.fit_image(image, 0.0, 0.0, "");
p.close_image(image);
```

The last parameter of `PDF_fit_image()` method is an option list which supports a variety of options for positioning, scaling, and rotating the image. Details regarding these options are discussed in Section 8.4, »Placing Images, Graphics, and imported PDF Pages«, page 190.

If an image file cannot be imported successfully `PDF_load_image()` returns an error code. If you need to know more details about the image failure, call `PDF_get_errmsg()` to retrieve a detailed error message.

Cookbook Code samples for image handling can be found in the `images` category of the *PDFlib Cookbook*.

Re-using image data. PDFlib supports an important PDF optimization technique for using repeated raster images. Consider a layout with a constant logo or background on multiple pages. In this situation it is possible to include the actual image data only once in the PDF, and generate only a reference on each of the pages where the image is used. Simply load the image file once, and call `PDF_fit_image()` every time you want to place the logo or background on a particular page. You can place the image on multiple pages, or use different scaling factors for different occurrences of the same image (as long as the image hasn't been closed). Depending on the image's size and the number of occurrences, this technique can result in enormous space savings.

Scaling and dpi calculations. PDFlib never changes the number of pixels in an imported image. Scaling either blows up or shrinks image pixels, but doesn't do any downsampling. A scaling factor of 1 results in a pixel size of 1 unit in user coordinates. In other words, the image is imported with its native resolution (or 72 dpi if it doesn't contain any resolution information) if the user coordinate system hasn't been scaled (since there are 72 default units to an inch).

Cookbook A full code sample can be found in the *Cookbook* topic `images/image_dimensions`. It shows how to get the dimensions of an image and how to place it with various sizes.

Color space of imported images. Except for adding or removing ICC profiles and applying spot or DeviceN color according to the options provided in `PDF_load_image()`, PDFlib generally preserves the native color space of an imported image. However, this is not possible for certain rare combinations.

PDFlib does not perform any conversion between RGB and CMYK. If such a conversion is required it must be applied to the image data before loading the image in PDFlib.

The sRGB ICC profile is assigned by default to all RGB images in BMP, JPG, JPX, PNG, or TIFF format which don't include an embedded ICC profile.

Multi-page images. PDFlib supports TIFF and JBIG2 images with more than one image, also known as multi-page image files. In order to use multi-page images use the `page` option in `PDF_load_image()`:

```
image = p.load_image("tiff", filename, "page=2");
```

The `page` option indicates that a multi-image file is to be used, and specifies the number of the image to use. The first image is numbered 1. This option may be increased until `PDF_load_image()` returns -1, signaling that no more images are available in the file.

Cookbook A full code sample for converting all images in a multi-image TIFF file to a multi-page PDF file can be found in the *Cookbook* topic `images/multi_page_tiff`.

Inline images. As opposed to reusable images, which are written to the PDF output as image XObjects, inline images are written directly into the respective content stream (page, pattern, template, or glyph description). This results in some space savings, but should only be used for small amounts of image data (up to 4 KB). The primary use of inline images is for bitmap glyph descriptions in Type 3 fonts; inline images are not recommended for other situations.

Inline images can be generated with `PDF_load_image()` and the `inline` option. Inline images cannot be reused, i.e., the corresponding handle must not be supplied to any call which accepts image handles. For this reason if the `inline` option has been provided `PDF_load_image()` internally performs the equivalent of the following code:

```
p.fit_image(image, 0, 0, "");  
p.close_image(image);
```

Inline images are only supported for `imagetype=jpeg` and `raw`. For other image types the `inline` option is silently ignored.

XMP metadata in images. Image files may contain XMP metadata. By default PDFlib ignores XMP metadata for images in the TIFF, JPEG, and JPEG 2000 image formats to reduce the output file size. However, the image metadata (if it contains valid XMP) can be attached to the generated Image XObject in the output PDF document with the following option of `PDF_load_image()`:

```
metadata={keepxmp=true}
```

8.1.2 Supported Image File Formats

PNG images. PDFlib supports all flavors of PNG images (ISO 15948). If a PNG image contains transparency information, the transparency is retained in the generated PDF (see Section 8.1.4, »Image Transparency«, page 165).

If a PNG image contains an sRGB chunk, the sRGB ICC profile is attached to the image unless the *honoriccprofile* option is *false* or another ICC profile has been assigned to the image with the *iccprofile* option. The rendering intent in the sRGB chunk will be used unless the *renderingintent* option has been supplied.

JPEG images. PDFlib supports the following flavors of JPEG images (ISO 10918-1):

- ▶ Grayscale, RGB (usually encoded as YCbCr, but direct RGB is also supported), and CMYK color
- ▶ Baseline and progressive JPEG compression.

These conditions cover all JPEG images in practical use. The restrictions above imply that images using some uncommon JPEG features cannot be imported, in particular arithmetic coding, lossless compression, bit depths other than 8 bits per component as well as non-standard features such as SmartScale and reversible RGB color transform. Note that the similar formats JPEG-LS (ISO 14495), sometimes called »lossless JPEG« and JPEG-XR (ISO 29199-2), formerly called HD Photo, are not supported.

JPEG images can be packaged in several file formats. PDFlib supports all common JPEG file formats and features:

- ▶ JFIF 1, which is generated by a wide variety of imaging applications.
- ▶ JPEG files written by Adobe Photoshop and other Adobe applications.
- ▶ PDFlib reads clipping paths from JPEG images created with Adobe Photoshop.
- ▶ PDFlib honors embedded ICC profiles in JPEG images unless the *honoriccprofile* option is set to *false*.
- ▶ If a JPEG image contains an Exif marker the color space information in the Exif marker is interpreted. If it indicates sRGB color space the sRGB ICC profile is attached to the image (unless the image contains an explicit embedded ICC profile, the *honoriccprofile* option is *false* or another ICC profile has been assigned to the image with the *iccprofile* option).
- ▶ The orientation entry in an Exif marker which specifies the desired image orientation is honored. It can be ignored (as many applications do) with the *ignore-orientation* option.

JPEG 2000 images. JPEG 2000 images (ISO 15444-2) require PDF 1.5 or above. PDFlib accepts JPEG 2000 images according to the following conditions:

- ▶ JP2 and JPX baseline images (usually **.jp2* or **.jpf*) are supported, subject to the color space conditions below. All color depth values in the range 1-38 bits are supported. The following color spaces are supported: sRGB, sRGB-grey, ROMM-RGB, sYCC, e-sRGB, e-sYCC, CIElab, ICC-based color spaces, and CMYK. PDFlib does not alter the original color space in the JPEG 2000 image file.
- ▶ (Unsupported) Raw JPEG 2000 code streams without JPX wrapper (often **.j2k*) with 1, 3, or 4 color components can be imported.
- ▶ Restricted or full ICC profiles embedded in the JPEG 2000 image file are kept, i.e. the *honoriccprofile* option is always *true*.

*Note JPM compound image files according to ISO 15444-6 (usually *.jpm) are not supported.*

Additional JPEG 2000 restrictions for PDF/X-4/5 (JPEG 2000 is not allowed in PDF/X-3 which is based on PDF 1.4):

- ▶ The number of color channels must be 1, 3, or 4.
- ▶ The bit depth of each color channel must be 1, 8, or 16.
- ▶ All color channels must have the same bit depth.
- ▶ Exactly one color space definition must be present in the JPEG 2000 image file.
- ▶ CMYK images can only be used if the output condition is a CMYK device or the *defaultcmyk* option has been supplied.

Additional JPEG 2000 restrictions for PDF/A-2 (JPEG 2000 is not allowed in PDF/A-1 which is based on PDF 1.4):

- ▶ The number of color channels must be 1, 3, or 4.
- ▶ All color channels must have the same bit depth.
- ▶ If the number of color space specifications in the JPEG2000 image file is larger than one, there must be exactly one color space specification that has the value 0x01 in the *APPROX* field.

JBIG2 images. PDFlib supports single- and multi-page flavors of JBIG2 images (ISO 14492). JBIG2 images always contain black/white pixel data.

Due to the nature of JBIG2 compression, several pages in a multi-page JBIG2 stream may refer to the same global segments. If more than one page of a multi-page JBIG2 stream is converted the global segments can be shared among the generated PDF images. Since the calls to *PDF_load_image()* are independent from each other you must inform PDFlib in advance that multiple pages from the same JBIG2 stream will be converted. This works as follows:

- ▶ When loading the first page all global segments are copied to the PDF. Use the following option list for *PDF_load_image()*:

```
page=1 copyglobals=all
```

- ▶ When loading subsequent pages from the same JBIG2 stream the image handle<N> for page 1 must be provided so that PDFlib can create references to the global segments which have already been copied with page 1. Use the following option list for *PDF_load_image()*:

```
page=2 imagehandle=<N>
```

The client application must make sure that the *copyglobals/imagehandle* mechanism is only applied to pages which are extracted from the same JBIG2 image stream. Without the *copyglobals* options PDFlib will automatically copy all required data for the current page.

TIFF images. PDFlib processes all relevant flavors of TIFF images:

- ▶ Compression schemes: uncompressed, CCITT (group 3, group 4, and RLE), ZIP (=Flate), PackBits (=RunLength), LZW, old-style and new-style JPEG, as well as some other rare compression schemes;
- ▶ Color space: black and white, grayscale, RGB, CMYK, CIElab, and YCbCr images; the color space in imported TIFF images is retained unmodified with the following exception: LZW-compressed TIFF images with CIElab color are converted to RGB, and do not retain the CIElab color space.

- ▶ Color depth must be 1, 2, 4, 8, or 16 bits per color component. 16-bit images require PDF 1.5.
- ▶ The BigTIFF format which extends the original TIFF format beyond 4GB.

The following TIFF features are processed when importing an image:

- ▶ TIFF files containing more than one image (see »Multi-page images«, page 160); use the *page* option to select a specific image within a TIFF file.
- ▶ Alpha channels or masks (see Section 8.1.4, »Image Transparency«, page 165) are honored unless the *ignoremask* option is set. You can explicitly select an alpha channel with the *alphachannelname* option.
- ▶ PDFlib honors clipping paths in TIFF images created with Adobe Photoshop and compatible programs unless the *ignoreclippingpath* option is set.
- ▶ PDFlib honors embedded ICC profiles in TIFF images unless the *honoriccprofile* option is set to *false*.
- ▶ If a TIFF image contains an Exif marker the color space information in the Exif marker is interpreted. If it indicates sRGB color space the sRGB ICC profile is attached to the image (unless the image contains an explicit embedded ICC profile, the *honoriccprofile* option is *false* or another ICC profile has been assigned to the image with the *iccprofile* option).
- ▶ The orientation tag which specifies the desired image orientation is honored. It can be ignored (as many applications do) with the *ignoreorientation* option.

BMP images. PDFlib supports the following flavors of BMP images:

- ▶ BMP versions 2 and 3;
- ▶ color depth 1, 4, and 8 bits per component, including 3 x 8 = 24 bit TrueColor. 16-bit images are treated as 5+5+5 plus one unused bit. 32-bit images are treated as 3 x 8 bit images (the remaining 8 bits are ignored).
- ▶ black and white or RGB color (indexed and direct color);
- ▶ uncompressed as well as 4-bit and 8-bit RLE compression;
- ▶ PDFlib does not mirror images if the pixels are stored in bottom-up order (this is a rarely used feature in BMP which is interpreted differently in applications).

Raw images. Uncompressed raw image data may be useful for special applications. Note that this format is unrelated to Camera Raw files created by Adobe applications.

In order to supply raw image data the *width*, *height* and *bpc* options must be supplied since PDFlib cannot derive the image dimensions from the image data. The color space can be derived from the *components* option: 1 component implies a grayscale image, 3 components an RGB image, and 4 components a CMYK image. Alternatively a spot or DeviceN color can be applied to the image via the *colorize* option. In this case the color space and number of color components are derived from the color space handle.

The length of the supplied image data must be equal to

$[\text{width} \times \text{components} \times \text{bpc} / 8] \times \text{height}$

bytes, with the bracketed term adjusted upwards to the next integer. Image samples are expected in top to bottom and left to right ordering (assuming no coordinate transformations have been applied). 16-bit samples must be provided with the most significant byte first (big-endian byte order). If *bpc* is smaller than 8, each pixel row begins on a byte boundary, and color values must be packed from left to right within a byte. Color channels are always interleaved, i.e. all color values for the first pixel must be supplied first,

followed by the color values for the second pixel, and so on. The user is responsible for supplying option values which match the image. Otherwise corrupt PDF output may be generated, and Acrobat may respond with the message *Insufficient data for an Image*.

The polarity of the color values is the same as for color-related options (see Chapter 4). Note that Adobe Photoshop uses the expected PDF polarity only for grayscale and RGB images, but inverse polarity for raw CMYK images and additional color channels: while PDF and PDFlib expect 0=no intensity, Photoshop assumes 0=maximum intensity. You can adjust the polarity by applying the *invert* option when loading the image.

8.1.3 Clipping Paths

PDFlib supports clipping paths in TIFF and JPEG images created with Adobe Photoshop or compatible programs. An image file may contain multiple named paths. Using the *clippingpathname* option of *PDF_load_image()* one of the named paths can be selected and will be used as a clipping path: only those parts of the image inside the clipping path will be visible; other parts remain invisible. This is useful to separate background and foreground, eliminate unwanted portions of an image, etc.

Alternatively, an image file may specify a default clipping path. If PDFlib finds a default clipping path in an image file it automatically applies it to the image (see figure 8.1). In order to prevent the default clipping path from being applied set the *honor-clippingpath* option in *PDF_load_image()* to *false*. If you have several instances of the same image and only some instances shall have the clipping path applied, you can supply the *ignoreclippingpath* option in *PDF_fit_image()* to disable the clipping path. When a clipping path is applied the bounding box of the clipped image is used as the basis for all calculations related to placing or fitting the image.

Cookbook A full code sample can be found in the *Cookbook* topic `images/integrated_clipping_path`. The *Cookbook* topic `path_objects/import_path_from_image` demonstrates how to visualize a clipping path.

The vector operations for describing the clipping path are written to the PDF output at each call to *PDF_fit_image()*. If an image with clipping path is placed in the document more than once it is strongly recommended to wrap the image within a template in order to reduce the output file size. This can be achieved with the *templateoptions* option of *PDF_load_image()*.



Fig. 8.1
Using a clipping path to separate foreground and background

8.1.4 Image Transparency

Image transparency is useful for a variety of artistic and other effects. For example, you can ignore irrelevant parts of an image and display only the interesting person or object. PDFlib supports several methods for image transparency:

- ▶ An alpha channel (also called soft mask) specifies transparency values for each image pixel. The alpha channel may be part of the imported image file or can be specified as a separate grayscale image. Soft masks with more than 1 bit per pixel are not allowed in PDF/A-1 and PDF/X-3.
- ▶ Chroma key masking specifies a single color value or a range of color value as transparent. The transparent color value(s) may come from the imported image file or can be specified via the *chromakey* option.
- ▶ Stencil masking uses a separate bitmap image which specifies transparent areas. The remaining image areas are painted with the current color or the main image.

Table 8.1 summarizes the available methods for masking an image with an alpha channel, chroma key value or stencil mask. The mask polarity differs for the alpha vs. stencil methods: the background shines through black areas of an alpha mask, but through white areas of a stencil mask.

Additional image color effects can be achieved with the *colorize* option (see Section 8.1.5, »Colorize Images with Spot or DeviceN Color«, page 168), blend modes and gstate soft masks (see Section 4.9, »Changing the Color of Objects«, page 90). For comparison the effects of colorizing an image with a spot or DeviceN color are also listed in Table 8.1.

Cookbook Full code samples can be found in the *Cookbook* topic `images/image_mask`.

Table 8.1 Supported methods for image transparency and colorizing images

<i>masking/colorizing method</i>	<i>base image</i>	<i>auxiliary mask image</i>	<i>effect of pixels in the mask or base image</i>
<i>alpha channel (soft mask) from image or masked option</i>			
<i>image with internal alpha channel</i>	<i>TIFF, PNG or JPEG 2000 image with internal alpha channel</i>	<i>(from image file)</i>	<i>(transparency according to the TIFF/PNG/JPEG 2000 image format)</i>
<i>image plus separate alpha channel image</i>	<i>any image; option <code>masked</code> refers to soft mask in auxiliary image</i>	<i>bitmap or grayscale image</i>	<i>black mask = background gray mask = semi transparent white mask = base image¹</i>
<i>chroma key masking from image or chromakey option</i>			
<i>image with internal chroma key value</i>	<i>PNG image with internal chroma key entry</i>	–	<i>chroma key color = transparent other image colors = base image</i>
<i>image plus separate chroma key value or range</i>	<i>any image except JPEG and JPEG 2000; option <code>chromakey</code> specifies transparent color value(s)</i>	–	<i>chroma key color(s) = transparent other image colors = base image</i>
<i>stencil masking with mask option</i>			
<i>paint through mask image as stencil</i>	–	<i>bitmap image loaded with option <code>mask</code></i>	<i>black mask = fill color white mask = background¹</i>

Table 8.1 Supported methods for image transparency and coloring images

masking/colorizing method	base image	auxiliary mask image	effect of pixels in the mask or base image
mask base image with stencil image	any image; option <i>masked</i> refers to stencil mask in auxiliary image	bitmap image loaded with option <i>mask</i>	black mask = base image white mask = background ¹
completely transparent image with constant opacity			
apply constant opacity to the full image	any image	–	all pixels appear with the same transparency
colorizing images with <i>colorize</i> option			
colorize image with spot color	grayscale image (>=1 bit per pixel); option <i>colorize</i> refers to spot color	–	black image = full spot color gray image = tinted spot color white image = white
colorize image with DeviceN color	n-channel image; option <i>colorize</i> refers to DeviceN color	–	0 in image = white intermediate = DeviceN tint 1 in image = full DeviceN color

1. The effect of the mask can be inverted by applying the *invert* option to the stencil or alpha image.

Images with an internal alpha channel. PDFlib reads an internal alpha channel from the following image formats (unless *ignoremask=true*):

- ▶ TIFF images may contain a single alpha channel which is used by PDFlib. Alternatively, a TIFF image may contain multiple alpha channels which are identified by name. If multiple channels are present in a TIFF image PDFlib by default uses the first alpha channel. You can explicitly select another channel by supplying its name with the *alphachannelname* option:

```
image = p.load_image("tiff", filename, "alphachannelname={apple}");
```

- ▶ PNG images may contain an alpha channel which is used by PDFlib.
- ▶ JPEG 2000 images may contain an alpha channel which is used by PDFlib.

Note As an alternative to a full alpha channel Photoshop can create transparent backgrounds in TIFF images using a proprietary format which is not supported in PDFlib. In order to use such transparent images with PDFlib you must save them in Photoshop in the TIFF file format and select Save Transparency in the TIFF options dialog box.

Using a separate grayscale image as alpha channel. As an alternative to an internal alpha channel in the image file you can use a second image as source of an alpha channel. All kinds of grayscale images are suitable for use as alpha channel. If the mask contains an embedded ICC profile this must be ignored with *honoriccprofile=false*. For TIFF images the *nopassthrough* option for *PDF_load_image()* is recommended to avoid multi-strip images. BMP images are oriented differently than other image types. For this reason BMP images must be mirrored along the x axis before they can be used as a mask.

White pixels in the alpha image result in the corresponding area of the base image being painted, while black mask pixels result in the background shining through. If the mask uses more than one bit per pixel, intermediate values blend the foreground image against the background, providing a transparency effect.

After loading the mask it is applied to the base image with the *masked* option:

```

// load the mask image which serves as alpha channel
mask = p.load_image("png", maskfilename, "");
if (mask == -1)
    throw new Exception("Error: " + p.get_errmsg());

// load base image and mask it
image = p.load_image(type, filename, masked=" + mask)
if (image == -1)
    throw new Exception("Error: " + p.get_errmsg());

p.fit_image(image, x, y, "");

```

Visualizing an image alpha channel. Sometimes you may want to display an image alpha channel as grayscale image. This may be useful for debugging or repurposing an alpha channel. It can be achieved by retrieving an image handle for the mask from the main image, using `PDF_info_image()` with the keyword `imagemask`:

```

image = p.load_image("auto", "image.tif", "errorpolicy=exception");

alpha = (int) p.info_image(image, "imagemask", "");
if (alpha != -1)
    p.fit_image(alpha, 0.0, 0.0, "");

```

The alpha channel is treated as grayscale image: transparent areas appear black.

Chroma key masking. Images may specify a single color or a range of colors to be masked out. Pixels with colors in the specified range are not painted so that the background shines through. In video technology this technique is called chroma key or blue screen masking. It is supported with the `chromakey` image option. This option expects n or $2n$ integers in the range 0 to $2^{\text{bitspercomponent}}-1$, where n is the number of components in the image color space.

Each pair in the list contains the inclusive lower and upper bounds of a color component range. Pixels where all color components (before applying any `decode` values or color inversion with the `invert` option) fall in the specified ranges are treated as transparent, i.e. they are not painted but allow the background to shine through.

If n values instead of n pairs are supplied each component range contains only a single color value, i.e. each list value describes both lower and upper bound for a color component.

Table 8.2 lists various examples of the `chromakey` option.

Table 8.2 Examples for the `chromakey` image option

<i>chromakey option</i>	<i>image color space</i>	<i>effect</i>
<code>chromakey={255 255 255}</code>	<i>RGB</i>	<i>treat white pixels as transparent</i>
<code>chromakey={0 255 128 255 0 255}</code>	<i>RGB</i>	<i>treat all pixels with more than 50% green as transparent</i>
<code>chromakey={242 255 242 255 242 255}</code>	<i>RGB</i>	<i>treat all colors lighter than 95% as transparent</i>
<code>chromakey={242 255 242 255 242 255}</code> <code>decode={0 0.95 0 0.95 0 0.95}</code>	<i>RGB</i>	<i>treat all colors lighter than 95% as transparent and spread the remaining colors to avoid sudden truncation at 95%</i>

Chroma key masking is applied automatically (unless *ignoremask=true*) for PNG images which contain a single transparent color value. If multiple color values with a corresponding alpha value are present, only the first one with an alpha value below 50 percent is used.

Stencil masks. Stencil masks are bitmap images with a bit depth of 1 where white pixels are treated as transparent: whatever content already exists on the page shines through the transparent parts of the image. The areas with 0 pixel values can be painted with the current fill color, or can be used to show parts of another image.

In order to colorize a stencil mask with the current fill color you must load the image with the *mask* option and set the current fill color before placing the image: Black pixels are colorized with the current fill color, while white areas remain unchanged:

```
mask = p.load_image("tiff", maskfilename, "mask");
if (mask == -1)
    throw new Exception("Error: " + p.get_errmsg());

p.set_graphics_option("fillcolor=red");
p.fit_image(mask, x, y, "");
```

If the stencil image is used to mask another image with the *masked* option, the pixels of that base image are visible in black areas of the stencil mask, and the background in white areas.

Applying constant opacity for a completely transparent image. You can make the complete image transparent (e.g. for use as background image) by applying constant opacity to all pixels when placing the image. This can be achieved with the *gstate* option and suboption *opacityfill*, e.g.

```
p.fit_image(mask, x, y, "gstate={opacityfill=0.5}");
```

8.1.5 Colorize Images with Spot or DeviceN Color

Similarly to stencil masks, where a color is applied to the non-transparent parts of a bitmap image PDFlib supports colorizing images with a spot or DeviceN color.

Colorize a grayscale image with a spot color. Black and white or grayscale images can be colorized with a spot color. To colorize an image with a spot color you must supply the *colorize* option when loading the image. This option contains a spot color handle created with *PDF_makespotcolor()*:

```
spot = p.makespotcolor("PANTONE Reflex Blue CV");

String optlist = "colorize=" + spot;
image = p.load_image("tiff", "image.tif", optlist);
```

As explained in Section 4.4, »Pantone, HKS, and custom Spot Colors«, page 79, spot color spaces usually expect *o=no color=white* in contrast to the grayscale color space where *o=black*. In order to ensure that white remains white PDFlib inverts the image polarity when colorizing an image with a spot color, i.e. dark image areas result in full spot color intensity. You can revert the color polarity with the *invert* image option.

Colorize an n-channel image with a DeviceN color. A DeviceN color handle created with `PDF_create_devicen()` can be supplied in the `colorize` image option to apply DeviceN color to an n-colorant image. In order to colorize an image with a DeviceN color the image data must contain the appropriate number N of color channels.

As explained in Section 4.5, »DeviceN Colors«, page 82, DeviceN color spaces expect `o=no color=white` in contrast to the grayscale color space where `o=black`. Unlike when colorizing an image with spot colors, PDFlib does *not* invert the image polarity when colorizing an image with a DeviceN color. As a result, colorizing with a spot color and colorizing with a DeviceN color with `N=1` expect different color polarities. You can invert the color polarity with the `invert` image option.

A code fragment for colorizing an image with a DeviceN color space can be found in »DeviceN color space based on CMYK process colors«, page 83.

Note Only raw images can be colorized with DeviceN color.

Cookbook A code sample can be found in the Cookbook topic `color/colorize_image_with_DeviceN_color`.

8.1.6 Modifying Color Values with a Decode Array

The color values of an image can be further modified with a linear decoding function. This method is supported with the `decode` image option. It expects `2xn` float or percentage values, where `n` is the number of components in the image color space. Applying decode arrays therefore requires knowledge of the image color space. Each pair of numbers in a decode array describes the target color values to which the lowest and highest component values `0` and `2bitspercomponent-1` are mapped; intermediate values are interpolated linearly. By default, image component values are mapped to the range `0..1` for most color spaces. Decode arrays can be used to apply certain color effects by compacting, spreading or clamping the color values. This technique may also be useful if the image is colorized or used as a mask.

Decode values have different effects on additive vs. subtractive color spaces. For example, increasing the values results in lighter RGB colors, but darker CMYK colors. For most color spaces (except Lab and Indexed) typical values are in the range `0..1`. Values outside this range are allowed, but the viewer clips the resulting color values to the appropriate range for the color space. This can be used for certain effects, e.g. to force a range of light colors to white.

Applying the `decode` option to an image which is used as a soft mask to another image can modify the effect of the mask, e.g. soften the mask.

The following image option shifts the black target values from the default range `0..1` to the new range `0.2..1.2`, i.e. it increases the black channel of a CMYK image by 20%, resulting in a stronger image:

```
decode={0 1 0 1 0 1 0.2 1.2}
```

Table 8.3 lists various examples of the `decode` option.

Table 8.3 Examples for the `decode` image option

<i>decode option</i>	<i>image color space</i>	<i>effect</i>
<code>decode={1 0 1 0 1 0}</code>	RGB	<i>invert image colors; same as invert</i>
<code>decode={-0.5 1.5 -0.5 1.5 -0.5 1.5}</code>	RGB	<i>increase contrast</i>

Table 8.3 Examples for the decode image option

<i>decode option</i>	<i>image color space</i>	<i>effect</i>
<code>decode={0.5 1.5 0.5 1.5 0.5 1.5}</code>	RGB	lighten all colors by 50%
<code>decode={-0.2 0.8 -0.2 0.8 -0.2 0.8}</code>	RGB	darken all colors by 20% (since 100%=white)
<code>decode={-0.2 0.8 -0.2 0.8 -0.2 0.8 -0.2 0.8}</code>	CMYK	lighten all colors by 20% (since 100%=dark)
<code>decode={0 1 0 1 0 1 0 0}</code>	CMYK	remove black channel, resulting in a duller image
<code>decode={0 1 0 1 0 1 0.2 1.2}</code>	CMYK	increase black channel by 20%, resulting in a stronger image
<code>decode={1 0 0 1 0 1 0 1}</code>	CMYK	invert Cyan channel while keeping other channels unchanged; unlikely to be useful except for artistic purposes

8.2 SVG Graphics

8.2.1 Supported SVG Flavors

PDFlib is a »Conforming High-Quality Static SVG Viewer« according to W3C nomenclature. PDFlib accepts SVG graphics as follows:

- ▶ PDFlib implements SVG 1.1 (Second Edition) as published by the W3C.
- ▶ The following Unicode formats and encodings are supported: UTF-8, UTF-16, ISO 8859-1 ... ISO 8859-15, ASCII
- ▶ CSS styling is available, but some CSS elements are unsupported.
- ▶ In addition to SVG files in plain text format Flate-compressed SVG files (*.svgz) are supported.
- ▶ Colors can be specified with additional color spaces according to the SVG Color 1.2 draft including a PDFlib-specific extension for the tint value of spot colors (see Section 8.2.6, »SVG Color Extension«, page 176).



See Section 8.2.8, »Unsupported SVG Features«, page 179, for restrictions.

8.2.2 SVG Processing Considerations

Basic SVG handling. Embedding vector graphics with PDFlib is easy to accomplish. First, the graphics file has to be opened with a PDFlib method which interprets the graphics and stores an internal representation in memory. The `PDF_load_graphics()` method returns a handle which serves as a graphics descriptor. This handle can be used in a call to `PDF_fit_graphics()`, along with positioning and scaling options:

```
graphics = p.load_graphics("auto", "graphics.svg", "");
if (graphics == -1)
    throw new Exception("Error: " + p.get_errmsg());

if (p.info_graphics(graphics, "fittingpossible", optlist) == 1)
    p.fit_graphics(graphics, 0.0, 0.0, "");
else
    System.err.println("Cannot place graphics: " + p.get_errmsg());
```

```
p.close_graphics(graphics);
```

The last parameter of `PDF_fit_graphics()` is an option list which supports a variety of options for positioning, scaling, and rotating. Details regarding these options are discussed in Section 8.4, »Placing Images, Graphics, and imported PDF Pages«, page 190.

Cookbook Code samples for SVG handling can be found in the `starter_svg` mini sample and the `graphics` category of the *PDFlib Cookbook*.

SVG content as inline graphics or template (Form XObject). PDFlib can create two different kinds of PDF output for imported SVG vector graphics:

- ▶ By default, PDFlib creates a separate template (Form XObject) for each imported SVG graphics which contains all required graphics operations. The template is created with default options; if custom options are required these can be supplied with the `templateoptions` option of `PDF_load_graphics()`. Templates have clearly defined transparency behavior. This is important if transparent elements of the SVG graphics interact with other elements in the same SVG or if existing page content should be partially visible through the SVG by applying opacity to the placed graphics. Templates also result in smaller PDF output if the graphics are placed more than once. As a disadvantage of the template method links within SVG graphics are not converted to PDF annotations.

Note the resulting SVG handle cannot be used as template handle although a template is created.

- ▶ If the `inline` option of `PDF_load_graphics()` is `true` the graphics data is written inline in the content stream of the page, pattern, template, or glyph description. The inline method is recommended if links within SVG graphics must be converted to PDF annotations (see option `convertlinks` of `PDF_fit_graphics()`). However, it suffers from several disadvantages: if the graphics contain overlapping transparent objects the output may contain wrong visual artifacts. Also, the opacity of the placed graphics cannot be changed reliably. If `PDF_fit_graphics()` is called more than once the graphics data is written again and again to the PDF output which increases output file size.

Using the same graphics in multiple PDF documents. Graphics can be loaded and closed independently from the current output document. When `PDF_load_graphics()` is called, an internal representation of the graphics is created. It is kept in memory until the corresponding call to `PDF_close_graphics()`. Keeping graphics in memory across documents has performance advantages in situations where the same graphics are placed in multiple output documents since the graphics must be loaded only once. For example, an application may load graphics with symbols, background artwork or company stationery once and call `PDF_fit_graphics()` in each document where the graphics are required.

Checking for SVG processing problems. While `PDF_load_graphics()` loads the SVG graphics, full processing and analysis is done only later in `PDF_fit_graphics()`, `PDF_close_graphics()` or `PDF_end_document()` depending on Form XObject creation and scope of loading. Since some error situations can be detected only during full processing, these functions could throw an exception if a problem was found (since they cannot return any error values). In order to avoid such an exception the graphics can be checked with the `fittingpossible` keyword of `PDF_info_graphics()`. It performs all process-

ing steps and doesn't create any output, but reports the success (or not) of SVG processing. If this check succeeds, `PDF_fit_graphics()` will not throw an exception when the image is placed. If an error occurs during the `fittingpossible` check, `PDF_info_graphics()` returns 0, regardless of `errorpolicy`. To summarize:

- ▶ The `fittingpossible` check prevents later exceptions in `PDF_fit_graphics()`, `PDF_close_graphics()` or `PDF_end_document()`. Since the PDF output would be unusable after an exception this is the recommended approach.
- ▶ Skipping the `fittingpossible` check speeds up SVG loading, but exceptions triggered by SVG data may occur later. This setting can be used to speed up SVG loading if exceptions in `PDF_fit_graphics()` are acceptable. For example, if the application converts a single SVG graphics file to a single PDF document without any additional page content this would be an acceptable approach.

The `fittingpossible` check uses the currently active global, document, and page options as well as the current output intent. It is therefore recommended to run this check only immediately before the actual call to `PDF_fit_graphics()`.

Network resources. SVG graphics may reference external resources such as fonts or images via URLs in the graphics file. By default, PDFlib retrieves such remote resources, i.e. the font or image is downloaded to the local machine. Details of network operation can be controlled with the `network` option of `PDF_set_option()`. The following protocols are supported: `file`, `http`, `https`, and `ftp`. In order to suppress all network operations use the suboption `disable` of the `network` option.

If it is known that multiple SVG graphics make use of the same network resources caching can be leveraged with the `downloadlifetime` option of `PDF_load_graphics()`

Note Network resources are not supported on IBM System i.

8.2.3 Size of SVG Graphics

SVG graphics specify the width and height in the `svg` element which defines the mapping of the SVG graphics to the target viewport (e.g. the browser window or some part of a PDF page). Often the size of the viewport is specified in absolute units, e.g.

```
<svg xmlns="http://www.w3.org/2000/svg" width="640mm" height="480mm">
```

PDFlib converts the `width` and `height` attributes to points and makes them available via the `graphicswidth` and `graphicsheight` keywords of `PDF_info_graphics()`. If the size is specified in pixels (`px`) PDFlib uses 1pt=1px. These values are also used to calculate the object box for fitting operations. The `svg` element may also contain the `viewBox` attribute which specifies a window inside the viewport.

You can override the size values specified in the SVG graphics file with the `forcedwidth` and `forcedheight` options. The suboption `clipping` of the `matchbox` option can be used to clip the graphics to some part of the object box.

SVG graphics without absolute size information. Some SVG graphics do not contain absolute size information since `width` and `height` are missing or contain only relative values such as in the following example:

```
<svg xmlns="http://www.w3.org/2000/svg" width="100%" height="100%">
```

This is used for SVG graphics which are not intended for standalone use. PDFlib supports this case if the user has specified a fitbox with the options *boxsize* and *fit-method=nofit*. The fitbox is used as object box. The SVG graphics will be clipped at the border of the fitbox. If no fitbox is specified PDFlib uses the *viewBox* attribute (if present) as object box.

If no absolute size information is available in the SVG graphics you can supply the *fallbackwidth* and *fallbackheight* options of *PDF_load_graphics()*. In the absence of these options PDFlib calculates the bounding box of the SVG graphics in the first call of *PDF_fit_graphics()* or *PDF_info_graphics()*. The calculated bounding box may be too small because line widths and oversized glyphs are not taken into account. In this case the calculated box can be enlarged with the *bboxexpand* option. By default the calculated bounding box is moved such that the graphics is located at the original position in the coordinate system (i.e. the box does not necessarily use the origin as corner).

8.2.4 Font Selection

Font selection algorithm. Font selection in SVG is controlled by the following properties:

font-family
font-style
font-weight

font-stretch
font-variant
font-size
font-size-adjust

Only the first three of these properties are relevant for selecting an external font. In order to select a suitable font, PDFlib constructs the following font names:

```
<font-family>,<font-weight>,<font-style>  
<font-family>-<font-weight><font-style>  
<font-family>,<font-normweight>,<font-style>  
<font-family>,<font-weight>,<font-normstyle>  
<font-family>,<font-normweight>,<font-normstyle>
```

where *<font-normweight>* is one of

Regular, Thin, Extralight, Light, Medium, Semibold, Bold, Extrabold, Black

and *<font-normstyle>* is

Italic

For example, with the following SVG font specification:

```
font-family="Tahoma" font-weight="Bold" font-style="Italic"
```

PDFlib will search for the font *Tahoma,Bold,Italic*, using the comma-separated PDFlib syntax for specifying font styles which can also be used to address Windows host fonts.

PDFlib then tries to load the font names listed above one after the other until the process is successful and a font could be loaded. The font names in this list can be used in font resource specifications, e.g.

```
p.set_option("FontOutline={<fontname>=<filename>}")
p.set_option("FontNameAlias={<fontname>=ArialMT}")
```

If all attempts fail, PDFlib tries to load the font with the name *<font-family>* and simulates the *Bold* and *Italic* properties if required.

Some browsers ignore the font selection properties if the specified font family cannot be found. Since PDFlib doesn't do this you must take care to make suitable fonts available via PDFlib's font configuration mechanism (see Section 6.3, »Loading Fonts«, page 116).

The *font-family* property may also contain multiple font family names, e.g.

```
font-family="Georgia, 'Minion Web', 'Times New Roman', Times, 'MS PMincho', serif"
```

In this case PDFlib tries to load the next font in the list if a particular font family could not be loaded. If some font for a *font-family* list could be loaded, PDFlib attempts to load the remaining *font-families* in the list as fallback fonts for the first loaded font (the master font, see Section 6.3.6, »Fallback Fonts«, page 127). If the master font already has fallback fonts because it has been loaded earlier, the new fallback fonts will be appended to the list of existing fallback fonts.

When PDFlib attempts to load a font for SVG graphics the following options are used by default:

```
embedding subsetting skipembedding={fstype latincore}
```

These options can be overridden with the option *defaultfontoptions* of *PDF_load_graphics()*.

Font configuration. On Windows systems PDFlib can access all fonts which are installed on the system (see Section 6.3.5, »Host Fonts on Windows and macOS«, page 126). For example, the SVG font specification

```
font-family="Verdana" font-weight="bold"
```

results in the PDFlib font name *Verdana,Bold*. On other operating systems PDFlib will find a font if a *FontOutline* resource has been specified in the following way:

```
<fontnamepattern>=<filename.xxx>
```

where *<fontnamepattern>* is one of the font name patterns above and *xxx* is the corresponding file name extension of the font outline file.

Suitable *FontOutline* resources with font names which match one or more of the font name patterns above can be created automatically with the *enumeratefonts* option of *PDF_set_option()*. Refer to Section 6.3.4, »Searching for Fonts«, page 122, for more information on font configuration.

Mapping generic SVG font families to PDF core fonts. PDFlib automatically maps the generic SVG font families *monospace*, *sans-serif* and *serif* to the Latin core fonts at the first occurrence of a generic SVG font family, using *FontNameAlias* resources of the following form:

```
p.set_option("FontNameAlias={monospace=Courier}")
p.set_option("FontNameAlias={monospace,Bold=Courier-Bold}")
```

The full list of generic font name mappings is as follows (there are no default mappings for the generic font families *curative* and *fantasy*):

monospace	Courier
monospace,Bold	Courier-Bold
monospace,Italic	Courier-Oblique
monospace,Bold,Italic	Courier-BoldOblique
sans	Helvetica
sans,Bold	Helvetica-Bold
sans,Italic	Helvetica-Oblique
sans,Bold,Italic	Helvetica-BoldOblique
sans-serif	Helvetica
sans-serif,Bold	Helvetica-Bold
sans-serif,Italic	Helvetica-Oblique
sans-serif,Bold,Italic	Helvetica-BoldOblique
serif	Times-Roman
serif,Bold	Times-Bold
serif,Italic	Times-Italic
serif,Bold,Italic	Times-BoldItalic

These mappings are executed only if no other suitable resource has been specified by the user before.

8.2.5 Dealing with missing Fonts and missing Glyphs

Missing fonts and the default font. If all font loading attempts fail, PDFlib attempts to load a default font which has the *font-family* name defined in the *defaultfontfamily* option of *PDF_load_graphics()*. By default, the *Arial Unicode MS* font is used if available, and Helvetica otherwise. It is highly recommended to either make the *Arial Unicode MS* font available in PDFlib's font configuration, or to specify another font with a large glyph complement in *defaultfontfamily*. For example, to configure the *NotoSans-Regular* font as last resort font use the following option:

```
defaultfontfamily={NotoSans-Regular}
```

Replacing individual fonts. In order to avoid a particular font which is not available or not desirable (e.g. because it does not contain enough glyphs, see below) you can map it to another font which is more suitable. Use the font name alias feature and *PDF_set_option()* for this purpose (see »Font name aliasing«, page 122, for details). For example, if you have Chinese text which is inappropriately set with the *Trebuchet MS* font which doesn't contain any Chinese glyphs, you can map it to *Arial Unicode MS* as follows:

```
p.set_option("FontnameAlias={ {Trebuchet MS}={Arial Unicode MS} }");
```

Keep in mind that font attributes are not automatically added. For example, if the *Trebuchet MS* font is used with the attribute *font-weight="bold"* you must create an alias for the bold version of the font:

```
p.set_option("FontnameAlias={ {Trebuchet MS,Bold}={Arial Unicode MS} }");
```

Visualizing missing glyphs. If a selected font does not contain a required glyph, the default replacement glyph is used instead, which means that no text will be visible at all. In order to visualize missing glyphs you can specify a visible replacement glyph using the *defaultfontoptions* option. For example, the following option for *PDF_load_graphics()* will display a question mark for all missing glyphs:

```
defaultfontoptions={replacementchar=?}
```

Specifying a particular fallback font for missing glyphs. If the font specified in SVG does not contain suitable glyphs for the text, no text will be visible. As a common example consider Chinese text which is attempted to be shown with a Western font which doesn't contain any Chinese glyphs. Of course the best solution would be to use suitable fonts in the SVG in the first place. However, if you have to deal with unsuitable fonts in SVG you can specify a fallback font in PDFlib. This fallback font is used whenever the original font does not provide a particular glyph.

The following option for *PDF_load_graphics()* calls for Arial Unicode MS as fallback font:

```
defaultfontoptions={fallbackfonts={{fontname={Arial Unicode MS}}}}
```

Note that the font specified in the *defaultfontfamily* option shown earlier is used when a font cannot be found, while the fallback technique applies when a font is available, but doesn't contain all required glyphs.

Specifying a global fallback font family. The *fallbackfontfamily* and *fallbackfontoptions* options can be used to specify a family of fallback fonts and corresponding options. While the *fallbackfonts* option within *defaultfontoptions* selects a single font for use as fallback font, *fallbackfontfamily* can be used to specify a family of fallback fonts, i.e. style attributes will be applied to this font family name, assuming that style variants of the specified font family are actually available. Example:

```
fallbackfontfamily={Arial}
```

8.2.6 SVG Color Extension

Cookbook A code samples for non-sRGB color in SVG can be found in the Cookbook topic *color/svg_color_extension*.

SVG by default uses the sRGB color space for text and vector graphics. PDFlib therefore treats SVG graphics as device-independent color for PDF/X and PDF/A. Embedded images in SVG may use other color spaces, though. For example, an embedded or referenced JPEG image may use the CMYK color space with or without an ICC profile. The default sRGB profile for RGB colors can be overridden with the *iccprofilergb* option of *PDF_load_graphics()*.

In addition to sRGB PDFlib supports additional color spaces using syntax extensions according to the SVG Color 1.2 draft specification. Table 8.4 lists supported SVG color spaces along with syntax examples. See Chapter 4, for general information about these color spaces and their treatment in PDFlib.

SVG syntax requires sRGB fallback colors for all variants. PDFlib uses the fallback color in case of syntax errors, missing spot color definitions and missing ICC profiles or if the option *forcesrgb* has been supplied.

You can colorize or decolorize SVG images with the help of blend modes (see Section 4.9.1, »Changing the Color with Blend Modes«, page 91).

ICC-based colors. The *icc-based* color specifications apply the specified ICC profile which is searched according to the *ICCprofile* resource category unless it references a local file. ICC profiles are ignored if the option *honoriccprofile=false* is supplied. Use the options *iccprofilegray/iccprofilergb/iccprofilecmyk* to override the profiles specified in the SVG graphics with other profiles.

Gray, RGB and CMYK device color spaces. The device-specific color spaces *device-gray*, *device-rgb* and *device-cmyk* are not recommended for general use as their rendering depends on the specific output device unless controlled by a PDF/A or PDF/X output intent ICC profile or suitable default color space. It is strongly recommended to work with ICC-based color spaces instead. Particular exceptions may be the use of a CMYK control strip to directly control tint values on a printing machine, or printing registration marks.

Device-dependent colors created by *device-gray/rgb/cmyk* attributes or referenced images can be mapped to ICC-based colors via the *iccprofile* suboption of *defaultimageoptions* or the *defaultgray/rgb/cmyk* suboptions of *templateoptions*.

Spot colors and DeviceN colors. The *icc-named-color* syntax can be used to specify spot colors. However, since named color ICC profiles are not supported the spot color must be known to PDFlib; otherwise the sRGB fallback color is used as alternate color. Custom spot colors must therefore be defined in your code (see »Custom spot colors«, page 81) before calling *PDF_load_graphics()*, e.g.

```
// Define custom spot color "CompanyRed" with Lab alternate values
p.set_graphics_option("fillcolor={spotname {CompanyRed} 1.0 {Lab 60 65 65}}");
```

The *device-nchannel* syntax can be used to select DeviceN colors which must be created with *PDF_create_devicen()* (see Section 4.5, »DeviceN Colors«, page 82) supplied to *PDF_load_graphics()* via the *devicencolors* option. This implies that only a single DeviceN color space can be used for each value of *N*.

Shadings. The color spaces listed in Table 8.4 can also be used in SVG gradients. However, PDF requires all stop colors for a gradient to be taken from the same color space. If a gradient uses stop colors from different color spaces the sRGB fallback colors for all stop colors are used.

Shadings created from spot colors retain the spot colors provided these have been created with Lab alternate values (which is true for all spot colors known internally to PDFlib). Spot colors created with other alternate color spaces such as CMYK are not retained, but result in sRGB shadings in the PDF output.

Note Since non-sRGB colors in SVG have not been fully standardized yet the syntax may change in the future.

Table 8.4 Extended color spaces in SVG (can be disabled with *forcesrgb*)

color space	SVG syntax examples	notes
-------------	---------------------	-------

Device-independent color spaces

Table 8.4 Extended color spaces in SVG (can be disabled with `forcesrgb`)

color space	SVG syntax examples	notes
<i>sRGB</i>	#FF0000	SVG 1.1 default color space
<i>ICC-based color</i>	#7F7F7F <code>icc-color(gray.icc, 0.5)</code> #CC6633 <code>icc-color(rgb.icc, 0.8, 0.4, 0.2)</code> #2B7FAB <code>icc-color(cmyk.icc, 0.8, 0.4, 0.2, 0.0)</code>	One, three or four color values must be supplied for Grayscale, RGB or CMYK profiles, respectively.
<i>CIELab</i>	#598237 <code>cielab(50, -25, 35)</code>	
<i>CIElch</i>	#FF007E <code>cielch(50, 127, 0)</code> #FF007E <code>cielchab(50, 127, 0)</code>	Lightness, chroma and hue (in degrees) are converted to CIELab.
<i>well-known spot color</i>	#FFB12D <code>icc-named-color(nmcl.icc, 'PANTONE 123 U')</code> #FFCE4C <code>icc-named-color(nmcl.icc, 'PANTONE 123 U', 0.5)</code> #994F5B <code>icc-named-color(nmcl.icc, 'HKS 18 Z')</code>	The profile name is ignored. As an extension to the SVG syntax PDFlib supports an optional tint value. If no tint value is present, 1.0 is assumed.
<i>custom spot color</i>	#FFE560 <code>icc-named-color(nmcl.icc, 'CompanyColor')</code> #FFF36D <code>icc-named-color(nmcl.icc, 'CompanyColor', 0.5)</code>	The sRGB fallback color is used as alternate color unless the spot color has already been defined earlier; the profile name is ignored. If no tint value is present, 1.0 is assumed.
<i>n-colorant</i>	#FFE560 <code>device-nchannel(0, 0.7, 0.2)</code> #FFF36D <code>device-nchannel(1 0 0 0.5 0.2 0)</code>	DeviceN color space handles must be supplied via the <code>devicenames</code> option. If no DeviceN color space handle for n has been supplied the sRGB fallback color is used.
Device-dependent color spaces		
<i>DeviceGray</i>	#7F7F7F <code>device-gray(0.5)</code>	PDF/A and PDF/X requirements for device color spaces must be obeyed.
<i>DeviceRGB</i>	#0000FF <code>device-rgb(0, 0, 255)</code>	PDF/A and PDF/X requirements for device color spaces must be obeyed.
<i>DeviceCMYK</i>	#00AEEF <code>device-cmyk(1, 0, 0, 0)</code>	PDF/A and PDF/X requirements for device color spaces must be obeyed.

8.2.7 SVG Contents beyond Vector Graphics and Text

Embedded images in SVG. PDFlib processes the *image* element in SVG and accepts all image formats discussed in Section 8.1, »Raster Images«, page 159, and nested SVG graphics. Image data may be embedded in the SVG file, located in an external file, or fetched from a remote location over the network.

Images in SVG graphics are processed automatically. However, in some situations you may want to supply certain image processing options. This can be achieved with the `defaultimageoptions` option of `PDF.load_graphics()`. For example, use the following option to apply antialiasing which improves the appearance of low-resolution images (this matches the SVG display of most browsers which apply anti-aliasing to images):

```
defaultimageoptions={interpolate}
```

If an image is not available (e.g. because the referenced external image file is missing) and the option `fallbackimage={}` (i.e. an empty option list) has been supplied, PDFlib creates a transparent gray checkerboard pattern. Suboptions of the option `fallbackimage` can be used to customize this pattern or to supply a custom image or template as fall-

back visualization. Use the following option if you want to avoid that missing images are silently replaced:

```
errorconditions={references={image}}
```

SVG links. Links in SVG graphics loaded with the *inline* option can be converted to interactive Link annotations in the generated PDF output. However, there are a few conditions which disable link creation (see PDFlib API Reference). The *contents* option of the PDF annotation is populated with the *xlink:title* attribute of the SVG link if present, and the target URI otherwise. In Tagged PDF mode a *Link* element and an associated OBJR element are created for the generated link unless the currently active item is a direct Artifact or pseudo element.

Conversion of SVG links to PDF links can be disabled with the *convertlinks* option of *PDF_fit_graphics()*.

Metadata in SVG. SVG graphics may contain XMP metadata. By default PDFlib ignores XMP metadata in SVG graphics to reduce the output file size. If a template is created from SVG (which is the default behavior) the metadata can be attached to the generated Form XObject with the following option of *PDF_load_graphics()*:

```
templateoptions={metadata={keepxmp=true}}
```

The contents of the *metadata*, *desc*, and *title* elements of an SVG graphic can be retrieved with *PDF_info_graphics()* according to the following pattern:

```
idx = (int) p.info_graphics(svg, "description", "");  
if (idx != -1)  
    description = p.get_string(idx, "");
```

8.2.8 Unsupported SVG Features

Treatment of unsupported features. By default, unsupported SVG features are ignored. As a result, output is created but some aspects of the graphics may be missing or wrong. This behavior can be changed with the *errorconditions* option of *PDF_load_graphics()*. Its suboptions specify conditions which trigger an error instead of being ignored. For example, with the following option list *PDF_load_graphics()* fails if the SVG graphics contain animated or scripted elements, e.g.:

```
errorconditions = {elements={animate script}}
```

Unsupported SVG elements. The following SVG elements are not supported and are ignored:

- ▶ elements for animation and scripting:

```
animate, animateColor, animateMotion, animateTransform, script, mpath, set
```

- ▶ elements for filters:

```
feBlend, feColorMatrix, feComponentTransfer, feComposite, feConvolveMatrix,  
feDiffuseLighting, feDisplacementMap, feDistantLight, feFlood, feFuncA, feFuncB,  
feFuncG, feFuncR, feGaussianBlur, feImage, feMerge, feMergeNode, feMorphology,  
feOffset, fePointLight, feSpecularLighting, feSpotLight, feTile, feTurbulence, filter
```

▶ other elements:

cursor, foreignObject, vkern

Restricted SVG elements, attributes and properties. The following attributes and properties are subject to restrictions:

- ▶ Some CSS directives are not supported.
- ▶ The font selection property *font-variant* is supported only with the keyword *small-caps*, and only for fonts containing the OpenType feature *smcp*.
- ▶ Combination of values of the text presentation property *text-decoration* are not supported. PDFlib doesn't draw the decoration elements as areas with separate fill and stroke color, but draws the decorations as lines. The lines are drawn with the fill color if present, otherwise with the stroke color.
- ▶ The attribute *rotate* for the *textPath* element is not supported.
- ▶ The property *unicode-bidi* is honored only for TrueType/OpenType fonts containing the required tables for Bidi text layout. PDFlib sets the options *shaping* and *script=_auto* in the option list of *PDF_fit_textline()*.
- ▶ The property *glyph-orientation-vertical* is supported according to SVG 1.1 except the angles 180° and 270° which are not supported.
- ▶ The attribute *preserveAspectRatio* for the *view* element is ignored.
- ▶ The property *overflow* for the *svg* element does not support the value *visible* if a FormXObject (template) is created, i.e. overflowing content is clipped. This can be avoided with the *inline* option.

Unsupported SVG properties. The following SVG properties are not supported and are ignored:

alignment-baseline, color-interpolation, color-interpolation-filters, color-rendering, cursor, dominant-baseline, enable-background, filter, flood-color, flood-opacity, glyph-orientation-horizontal, image-rendering, lighting-color, pointer-events, shape-rendering, text-rendering

Unsupported attributes of supported SVG elements. The following attributes of supported SVG elements are not supported and are ignored:

baseProfile (svg)
contentScriptType (svg)
contentStyleType (svg)
externalResourcesRequired (all elements)
method (textPath)
on* (all elements)
spacing (textPath)
textLength and lengthAdjust for vertical text (text, textpath, tref, tspan)
version (svg)
zoomAndPan (svg)
xlink:role (all elements)
xlink:show (all elements)
xlink:type (all elements)

8.3 Importing PDF Pages with PDI

Note All functions described in this section require PDFlib+PDI or the PDFlib Personalization Server PPS (which includes PDI). The PDF import library (PDI) is not contained in the PDFlib base prod-

uct. Although PDI is included in all binary editions of PDFlib, a license key for PDFlib+PDI or PPS is required to use it.

8.3.1 PDI Features and Applications

If PDI (PDF import library) is available, pages from existing PDF documents can be imported. PDI prepares pages from existing PDF documents for use with PDFlib. Conceptually, imported PDF pages are treated similarly to imported raster images: you open a PDF document, choose a page to import, and place it on an output page. You can apply any of PDFlib's transformation functions for translating, scaling, rotating, or skewing the imported page. Imported pages can be combined with new content by using PDFlib's text or graphics functions after placing the imported PDF page on the output page (think of the imported page as the background for new content). Using PDFlib and PDI you can easily accomplish the following tasks:

- ▶ overlay two or more pages from multiple PDF documents (e.g., add stationary to existing documents in order to simulate preprinted paper stock);
- ▶ place PDF ads in existing documents;
- ▶ clip the visible area of a PDF page in order to get rid of unwanted elements (e.g., crop marks), or scale pages;
- ▶ impose multiple pages on a single sheet for printing;
- ▶ process multiple PDF/X or PDF/A documents to create a new PDF/X or PDF/A file;
- ▶ copy the PDF/X or PDF/A output intent of a file;
- ▶ add some text (e.g., headers, footers, stamps, page numbers) or images (e.g., company logo) to existing PDF pages;
- ▶ copy all pages from an input document to the output document, and place barcodes on the pages;
- ▶ use the pCOS interface to query arbitrary properties of a PDF document (see pCOS Path Reference for details).

In order to place a PDF background page and populate it with dynamic data (e.g., mail merge, personalized PDF documents on the Web, form filling) we recommend using PDI along with PDFlib blocks (see Chapter 13).

8.3.2 Using PDFlib+PDI

Cookbook Code samples regarding PDF import issues can be found in the `pdf_import` category of the *PDFlib Cookbook*.

General considerations. The following items can be imported from PDF documents in addition to the page contents:

- ▶ Annotations (see »Importing annotations«, page 183).
- ▶ Form fields (see »Importing form fields«, page 184).
- ▶ Actions (see »Importing actions and JavaScript«, page 184).
- ▶ Structure element tags (see »Importing structure information from Tagged PDF documents«, page 185, for details).
- ▶ Layer definitions (see »Importing pages with layers«, page 185).
- ▶ (PPS only) PDFlib Blocks can be imported with `PDF_process_pdi()` and the option `action=copyallblocks` or `copyblock` (see Section 13.9.2, »Importing PDFlib Blocks«, page 373).

You cannot re-use individual resources of imported pages with other PDFlib functions. For example, re-using fonts from imported documents for some other content is not possible. Instead, all required fonts must be configured in PDFlib. If multiple imported documents contain embedded font data for the same font, PDI does not remove duplicate font data by default; this can be achieved with the *optimize* document option. On the other hand, if fonts are missing from an imported PDF they will also be missing from the generated PDF output file. As an optimization you should keep the imported document open as long as possible in order to avoid the same fonts to be embedded multiple times in the output document.

PDFlib+PDI uses the template feature (Form XObjects) for placing imported PDF pages on the output page.

Code fragment for importing PDF pages. Dealing with pages from existing PDF documents is possible with a very simple code structure. The following code snippet opens a page from an existing document and copies the page contents to a new page in the output PDF document (which must have been opened before):

```
int    doc, page, pagecount, pageno = 1;
String filename = "input.pdf";

if (p.begin_document(outfilename, "") == -1) {...}
...

doc = p.open_pdi_document(infile, "");
if (doc == -1)
    throw new Exception("Error: " + p.get_errmsg());

/* Number of pages in the input document; useful for importing all pages */
pagecount = (int) p.pcos_get_number(doc, "length:pages");

page = p.open_pdi_page(doc, pageno, "");
if (page == -1)
    throw new Exception("Error: " + p.get_errmsg());

/* dummy page size, will be modified by the adjustpage option */
p.begin_page_ext(20, 20, "");
p.fit_pdi_page(page, 0, 0, "adjustpage");
p.close_pdi_page(page);
...add more contents to the page using PDFlib functions...
p.end_page_ext("");
p.close_pdi_document(doc);
```

The last parameter to *PDF_fit_pdi_page()* is an option list which supports a variety of options for positioning, scaling, and rotating the imported page. Details regarding these options are discussed in Section 8.4, »Placing Images, Graphics, and imported PDF Pages«, page 190.

8.3.3 Document and Page-related Checks

Document-related checks. PDFlib+PDI processes all kinds of PDF documents regardless of PDF version number or features used within the file. PDFlib+PDI implements a repair mode for damaged PDFs so that even certain damaged documents can be opened.

However, in rare cases a PDF document or a particular page of a document may be rejected by PDI.

If a PDF document or page can't be imported successfully `PDF_open_pdi_document()` and `PDF_open_pdi_page()` return an error code. If you need to know more details about the failure you can query the reason with `PDF_get_errmsg()`. Alternatively, you can set the `errorpolicy` option to `exception`, which will result in an exception if the document cannot be opened.

Page-related checks. The following checks are done in `PDF_open_pdi_page()`:

- ▶ Pages from PDF documents which use a higher PDF version number than the PDF output document that is currently being generated cannot be imported. The reason is that PDFlib could no longer ensure that the output conforms to the requested PDF version after a PDF with a higher version number has been imported. Solution: set the version of the output PDF to the required level using the `compatibility` option in `PDF_begin_document()`.

`PDF 1.7ext3` and `PDF 1.7ext8` documents are compatible with PDF 1.7 as far as PDI is concerned.

In PDF/A mode the input PDF version number is not relevant since the PDF version header must be ignored in PDF/A.

If a document uses a higher PDF version header although it is known to conform to an older PDF version you can use the `ignorepdfversion` option of `PDF_open_pdi_document()`.

- ▶ PDF/A, PDF/X, PDF/VT or PDF/UA documents which are incompatible to the corresponding PDF/A, PDF/X, PDF/VT or PDF/UA status of the current output document are rejected. See the following sections for more information:
 - ▶ Section 12.2.7, »Importing PDF/A Documents with PDI«, page 312;
 - ▶ Section 12.3.6, »Importing PDF/X Documents with PDI«, page 322;
 - ▶ Section 12.4.7, »Importing PDF/X and PDF/VT Documents with PDI«, page 330;
- ▶ Section 11.4.4, »Importing Pages from PDF/UA Documents with PDI«, page 298.

8.3.4 Importing Annotations, Form Fields and Actions

Importing annotations. By default, PDI imports the page contents and all annotations which may be present on the imported page. You can supply the `useannots` option of `PDF_fit_pdi_page()` to import only selected annotation types such as *Link* annotations. Actions which are associated with annotations are also imported. Link destinations in the document are adjusted according to the new placement of pages. For example, intra-document links still work even if you change the order of imported pages. Link destinations are recalculated, e.g. if the target page has been scaled or rotated. However, some interactive elements cannot be reproduced successfully depending on the contents and structure of the generated output document. For example, if you import a page containing a link to the next page, but the target page is not imported – where should the link point to? In general PDFlib tries to make as many interactive elements as possible functional in the new document, but this cannot be guaranteed in all cases.

Annotation import works only in *page* scope (e.g. not when a PDF page is placed on a template) and is subject to the following restrictions:

- ▶ If a link target is not available (e.g. because the target page is not imported) a link without any action is created.

- ▶ Named destinations are resolved, but the names are not preserved. This means that links jump to the target page, but the symbolic destination names are gone.
- ▶ Importing *PrinterMark* and *Watermark* annotations is not supported.
- ▶ Some properties in annotation dictionaries are not imported, including the Movie activation keys *Start*, *Duration*, *Rate*, *Synchronous*.
- ▶ Importing PDF 2.0 structure destinations is not supported.

If a page containing one of the annotation types *Caret*, *FreeText*, *Line* or *Stamp* is imported and the annotation doesn't have an appearance stream, access to the Noto fonts must be configured. Both fonts are included in the PDFlib distribution (see Section 1.3, »Configuration Requirements«, page 17).

Importing form fields. By default, PDI imports only the page contents and ignores form fields. You can supply the *usefields* option of *PDF_fit_pdi_page()* to import all fields or only selected types such as text fields. Form fields can only be imported from a single PDF input document, but not from multiple documents. An imported page can be placed with form fields only once. Also, form fields can be created either by importing from a single PDF document or with *PDF_create_field()*. Both methods cannot be combined. However, imported form fields can be filled with *PDF_create_field()*.

Form field import works only in *page* scope (e.g. not when a PDF page is placed on a template) and is subject to the following restrictions:

- ▶ Location and size of form field rectangles are adjusted to the placement of the page, e.g. scaled up or down.
- ▶ Cryptographic signature data of signature fields is cleared; in this case the visual representation is also recreated. Invisible signatures are ignored.

Importing actions and JavaScript. PDF documents may contain actions which are activated by various trigger events. Import of actions is organized as follows:

- ▶ Actions which are located at annotations and form fields are automatically imported along with the corresponding annotation or form field.
- ▶ Importing the following action types is not supported:
GoToE, *GoToDp*, *Thread*, *Sound*, and *Trans*.
- ▶ Document actions can be imported with the *useactions* option of *PDF_open_pdi_document()*.
- ▶ Page actions can be imported with the *useactions* option of *PDF_fit_pdi_page()*.
- ▶ JavaScript actions for annotations, form fields, pages or the document may rely on JavaScript routines stored on the document level. These scripts can be imported with the *usejavascript* option of *PDF_open_pdi_document()*. Setting this option to *true* is recommended if fields are imported which rely on document-level JavaScript. Document-level JavaScript can also be created with the *script* and *scriptname* options of *PDF_create_action()*.

8.3.5 Additional Aspects of imported PDF Documents

Dimensions of imported PDF pages. Imported PDF pages are handled similarly to imported raster images, and can be placed on the output page using *PDF_fit_pdi_page()*. By default, PDI imports the page exactly as it is displayed in Acrobat, in particular:

- ▶ cropping is retained (in technical terms: if a *CropBox* is present, PDI favors the *CropBox* over the *MediaBox*; see Section 3.2.2, »Page Size«, page 61);

- ▶ rotation which has been applied to the page is retained.

The *cloneboxes* option instructs PDFlib+PDI to copy all page boxes of the imported page to the generated output page, effectively cloning all page size aspects.

Alternatively, you can use the *pdiusebox* option to explicitly instruct PDI to use any of the MediaBox, CropBox, BleedBox, TrimBox or ArtBox entries of a page (if present) for determining the size of the imported page.

Finally, you can use the *transform* option of *PDF_open_pdi_page()* to apply arbitrary transformations to the imported page, e.g. scale it up or down or rotate it.

Color handling. PDFlib+PDI does not change the color of imported PDF documents in any way. For example, if a PDF contains ICC color profiles these are retained in the output document. If a page contains a transparency group entry it will be copied to the generated Form XObject unless the *transparencygroup* option of *PDF_open_pdi_page()* asks for different treatment. By default a suitable transparency group is created for the imported page if none is present on the imported page, taking into account all applicable PDF/A, PDF/X, and PDF/VT requirements.

Importing structure information from Tagged PDF documents. By default, tags are imported if both input and output document are tagged. However, tag import can be disabled with the *usetags* options of *PDF_open_pdi_document()* and *PDF_open_pdi_page()*. For more information refer to Section 11.3, »Importing Tagged PDF Pages with PDI«, page 290.

Importing pages with layers. PDI always imports the contents of all layers (technically known as *optional content*) on a page. The layer definitions including the layers' visibility state are also imported provided the layer is used on any of the imported pages. However, import of layer definitions can be disabled with the *uselayers* options of *PDF_open_pdi_document()*. In order to work with *uselayers=false* the generated document must not contain any layers at all, i.e. all PDF documents with layers must be opened with *uselayers=false* and *PDF_define_layer()* must not be called.

For more control for arranging the imported layers you can use the *parenttitle* of *PDF_open_pdi_document()* which creates a hierarchical title layer in the layer list on top of the imported layers (e.g. for supplying a file name). The *parentlayer* option works similarly but expects a handle to a user-defined layer.

Importing georeferenced PDF. When importing georeferenced PDF with PDI the geospatial information is kept if it has been created with one of the following methods (image-based geospatial reference):

- ▶ with PDFlib and the *georeference* option of *PDF_load_image()*
- ▶ by importing an image with geospatial information in Acrobat.

The geospatial information is lost after importing a page if it has been created with one of the following methods (page-based geospatial reference):

- ▶ with PDFlib and the *viewports* option of *PDF_begin/end_page_ext()*
- ▶ by manually geo-registering a PDF page in Acrobat.

Optimization across multiple imported documents. While PDFlib itself creates highly optimized PDF output, imported PDF may contain redundant data structures which in some cases can be optimized. In addition, importing multiple PDFs may bloat the out-

put file size if multiple files contain identical resources, e.g. fonts. In this situation you can use the *optimize* option of `PDF_begin_document()`. It detects redundant objects in imported files and removes them without affecting the visual appearance or quality of the generated output.

Encrypted PDF documents and the »shrug« feature. In order to import pages from encrypted documents (i.e., files with permission settings or password) the corresponding master password must be supplied. Encrypted PDF documents without the master password are rejected by default. However, they can be opened for querying information with pCOS (as opposed to importing pages) by setting the *infomode* option of `PDF_open_pdi_document()` to *true* (exception to the *infomode* rule: documents created with the Distiller setting *Object Level Compression: Maximum* cannot be opened even in info mode).

With the *shrug* feature pages from protected documents can be imported without master password, assuming the user accepts responsibility for respecting the document author's rights. By using the *shrug* feature the user asserts that he or she does not violate any document authors' rights. PDFlib GmbH's terms and conditions require that users respect document author's rights.

If all of the following conditions are true, the *shrug* feature is enabled:

- ▶ The *shrug* option has been supplied to `PDF_open_pdi_document()`.
- ▶ The document requires a master password but it has not been supplied to `PDF_open_pdi_document()`.
- ▶ If the document requires a user (open) password, it must have been supplied to `PDF_open_pdi_document()`.

The *shrug* feature has the following effects:

- ▶ Pages can be imported although the master password has not been supplied.
- ▶ The pCOS pseudo object *shrug* is set to *true/1*.
- ▶ pCOS runs in full mode (instead of restricted mode), i.e. the *pcosmodename* pseudo object is set to the value *full*.

8.3.6 Cloning Document-related Properties

PDF documents have a variety of properties and may include many features beyond the page contents. When importing PDF documents with PDFlib+PDI you have detailed control over which aspects of an imported document should be replicated in the generated output document. In some cases a simple option instructs PDFlib+PDI to clone an aspect, while in other cases you can use the pCOS interface to retrieve some data from the imported document and then feed this data to the corresponding PDFlib API methods and options.

Table 8.5 Cloning document-related aspects of imported PDF documents

PDF item	pCOS path	clone or create with PDFlib
<i>The following properties can be cloned with options:</i>		
<i>document output intent</i>	/Root/OutputIntents	<code>PDF_process_pdi()</code> : option <i>action=copyoutput-intent</i>
<i>document-level actions</i>	<i>destpage</i> , /Root/OpenAction, /Root/AA	<code>PDF_open_pdi_document()</code> : option <i>useactions</i>

Table 8.5 Cloning document-related aspects of imported PDF documents

PDF item	pCOS path	clone or create with PDFlib
JavaScript name tree	names/JavaScript	<i>PDF_open_pdi_document(): option usejavascript</i>
named destinations	names/Dests	Named destinations used by annotation or field actions are resolved, but the names are not imported. <i>PDF_add_nameddest()</i>
base URI	/Root/URI	The Base URI is automatically applied to imported Link actions. <i>PDF_begin/end_document(): option uri</i>
The following properties can be cloned by retrieving objects with pCOS and feeding them to PDFlib:		
document attachments	names/EmbeddedFiles	<i>PDF_begin/end_document(): option attachments</i>
Tagged PDF status	tagged	<i>PDF_begin_document(): option tagged</i>
page layout when document is opened	/Root/PageLayout	<i>PDF_begin/end_document(): option pagelayout</i>
document open mode	/Root/PageMode	<i>PDF_begin/end_document(): option openmode</i>
viewer preferences	/Root/ViewerPreferences	<i>PDF_begin/end_document(): option viewer-preferences</i>
document language	/Root/Lang	<i>PDF_end_document(): option lang or tagging option lang</i>
encryption	encrypt/algorithm encrypt/master encrypt/user encrypt/attachment encrypt/plainmetadata encrypt/noaccessible etc.	<i>PDF_begin_document(): options attachmentpassword, masterpassword, permissions, userpassword</i>
document-level XMP	/Root/Metadata	<i>PDF_begin/end_document(): option metadata</i>
document info fields	/Info	<i>PDF_set_info()</i>
linearization	linearized	<i>PDF_begin_document(): option linearize</i>
PDF/A, PDF/UA, PDF/X, PDF/VT standard status	pdfa, pdfua, pdfx, pdfvt	<i>PDF_begin_document(): options pdfa, pdfua, pdfx, pdfvt</i>
PDF version	pdfversionstring	<i>PDF_begin_document(): option compatibility</i>
associated files on document level	/Root/AF	<i>PDF_end_document(): option associatedfiles</i>
associated files on the structure hierarchy	/Root/StructTreeRoot/AF	<i>PDF_end_document(): option structureassociatedfiles</i>
The following properties cannot be cloned easily:		
bookmarks	bookmarks[...]	<i>PDF_create_bookmark()</i>
DPart and Document Part Metadata (DPM)	/Root/DPartRoot	<i>PDF_begin/end_dpart(): option dpm</i>
portfolio	/Root/Collection	<i>PDF_begin/end_document(): option portfolio, PDF_add_portfolio_file(), PDF_add_portfolio_fol-der()</i>

Table 8.5 Cloning document-related aspects of imported PDF documents

PDF item	pCOS path	clone or create with PDFlib
<i>digital signature and timestamp</i>	<code>signaturefields[...]</code>	<i>Cloning signatures doesn't make sense since they would be invalid. New signatures and timestamps can be created with PDFlib PLOP DS.</i>
<i>usage rights for enabling editing tools in Reader</i>	<code>usagerights</code>	<i>Usage rights can only be created with Adobe tools due to company-specific digital signatures.</i>
<i>class map for Tagged PDF</i>	<code>/Root/StructTreeRoot/ClassMap</code>	<i>unsupported</i>
<i>article threads and beads</i>	<code>articles</code>	<i>unsupported</i>
<i>document-level PieceInfo</i>	<code>/Root/PieceInfo</code>	<i>unsupported</i>

8.3.7 Cloning Page-related Properties

Table 8.6 Cloning page-related aspects of imported PDF documents

PDF item	pCOS path	clone or create with PDFlib
The following properties can be cloned with options:		
page content	pages[...]	<i>PDF_fit_pdi_page()</i>
page size	pages[...]/MediaBox and pages[...]/CropBox	<i>PDF_fit_pdi_page()</i> : option adjustpage
page geometry (boxes and rotation)	pages[...]/MediaBox etc. pages[...]/Rotate	<i>PDF_open_pdi_page()</i> and <i>PDF_fit_pdi_page()</i> : option cloneboxes
annotations	pages[...]/annots	<i>PDF_fit_pdi_page()</i> : option useannots
form fields	pages[...]/fields	<i>PDF_fit_pdi_page()</i> : option usefields
page open and close actions	pages[...]/AA/O and pages[...]/AA/C	<i>PDF_fit_pdi_page()</i> : option useactions
Blocks for PDFlib Personalization Server (PPS)	pages[...]/blocks	<i>PDF_process_pdi()</i> : options action=copyblock and action=copyallblocks
layers	/Root/OCProperties/OCGs	<i>PDF_open_pdi_document()</i> : option uselayers
structure element tags for page content	<i>PDF_info_pdi_page()</i> : keyword topleveltag	<i>PDF_open_pdi_document()</i> and <i>PDF_open_pdi_page()</i> : option usetags
page output intent	pages[...]/OutputIntents	<i>PDF_process_pdi()</i> : option action=copypage-outputintent
The following properties can be cloned by retrieving objects with pCOS and feeding them to PDFlib:		
page-level XMP	pages[...]/Metadata	<i>PDF_begin/end_page_ext()</i> : option metadata
tab order for annotations and form fields	pages[...]/Tabs	<i>PDF_begin/end_page_ext()</i> : option taborder
user units	pages[...]/UserUnit	<i>PDF_begin/end_page_ext()</i> : option userunit
viewport	pages[...]/VP	<i>PDF_begin/end_page_ext()</i> : option viewports
page transition	pages[...]/Trans	<i>PDF_begin/end_page_ext()</i> : options duration and transition
associated files on page level	pages[...]/AF	<i>PDF_begin/end_page_ext()</i> : option associatedfiles
associated files on image level	pages[...]/images[...]/AF	<i>PDF_load_image()</i> : option associatedfiles
associated files on Form XObject level	pages[...]/templates[...]/AF	<i>PDF_load_graphics()</i> , <i>PDF_open_pdi_page()</i> , <i>PDF_begin_template_ext()</i> : option associatedfiles
associated files on annotation XObject level	pages[...]/annots[...]/AF	<i>PDF_create_annotation()</i> : option associatedfiles
page labels	pages[...]/label	<i>PDF_begin_page_ext()</i> : option label
The following properties cannot be cloned:		
page-level PieceInfo	pages[...]/PieceInfo	unsupported
presentation steps	pages[...]/PresSteps	unsupported

8.4 Placing Images, Graphics, and imported PDF Pages

The functions `PDF_fit_image()` for placing raster images and templates, `PDF_fit_graphics()` for placing graphics and `PDF_fit_pdi_page()` for placing imported PDF pages offer a wealth of options for controlling the placement on the page. This section demonstrates the most important options by looking at some common application tasks. A complete list and descriptions of all options can be found in the PDFlib API Reference.

All samples in this section work the same for raster images, templates, graphics and imported PDF pages. Although code samples are only presented for raster images we talk about placing objects in general. Before calling any of the *fit* functions a call to `PDF_load_image()`, `PDF_load_graphics()` or `PDF_open_pdi_document()` and `PDF_open_pdi_page()` must be issued. For the sake of simplicity these calls are not repeated here.

Cookbook Code samples regarding images, graphics and imported PDF pages can be found in the `images`, `graphics`, and `pdf_import` categories of the PDFlib Cookbook.

8.4.1 Simple Object Placement

Positioning an image at the reference point. By default, an object will be placed in its original size with the lower left corner at the reference point. In this example we will place an image with the bottom centered at the reference point (*o*, *o*):

```
p.fit_image(image, 0, 0, "position={center bottom}");
```

Similarly, you can use the *position* option with another combination of the keywords *left*, *right*, *center*, *top*, and *bottom* to place the object at the reference point.

Placing an image with scaling. The following variation will place the image while modifying its size:

```
p.fit_image(image, 0, 0, "scale=0.5");
```

This code fragment places the object with its lower left corner at the point (*o*, *o*) in the user coordinate system. In addition, the object will be scaled in *x* and *y* direction by a scaling factor of 0.5, which makes it appear at 50 percent of its original size.

Cookbook A full code sample can be found in the Cookbook topic `images/starter_image`.

8.4.2 Placing an Object at a Point or Line or in a Box

In order to position an object an additional box with specified width and height can be used. The gray box or line in the figures below is depicted for visualizing the box size only; it is not part of the actual output.

Placing an object in a box doesn't make sense for *fitmethod=nofit* since the object is only positioned in this case, but not scaled. The *boxsize* option can be used to specify a horizontal line, vertical line or real box for object placement:

```
boxsize={100 0}           horizontal line
boxsize={0 100}          vertical line
boxsize={100 200}        box
```

In the examples below we will fit the object into the box with various fitting methods.

Automatic fitting in a box. With *fitmethod=auto* PDFlib scales the image to fit into the box without distorting it: if it fits into the box no scaling is applied. Otherwise the size is reduced while preserving the aspect ratio of width and height. figure 8.2a, figure 8.2b, and figure 8.2c demonstrate how PDFlib reduces the image size as the size of the fitbox decreases from initially *boxsize={70 45}* to *boxsize={70 30}* and further to *boxsize={30 30}*.

Fig. 8.2 Fitting an image into a box subject to various fit methods

Generated output	Option list for <i>PDF_fit_image()</i>
a) 	<code>boxsize={70 45} position=center fitmethod=auto</code> (no scaling necessary)
b) 	<code>boxsize={70 30} position=center fitmethod=auto</code> (scaled down to adjust for smaller box height)
c) 	<code>boxsize={30 30} position=center fitmethod=auto</code> (scaled down to adjust for smaller box height and width)
d) 	<code>boxsize={70 45} position=center fitmethod=meet</code>
e) 	<code>boxsize={35 45} position=center fitmethod=meet</code>
f) 	<code>boxsize={70 45} position=center fitmethod=entire</code>
g) 	<code>boxsize={30 30} position=center fitmethod=clip</code>
h) 	<code>boxsize={30 30} position={right top} fitmethod=clip</code>

Fitting an image in the center of a box. In order to center an image within a pre-defined rectangle you don't have to do any calculations, but can achieve this with suitable options. With *position=center* we place the image in the center of the box, 70 units wide and 45 high (*boxsize={70 45}*). Using *fitmethod=meet*, the image is proportionally resized until its height completely fits into the box (see figure 8.2d).

Decreasing the box width from 70 to 35 units forces PDFlib to scale down the image until its width completely fits into the box (see figure 8.2e).

Fig. 8.3
The rotate option

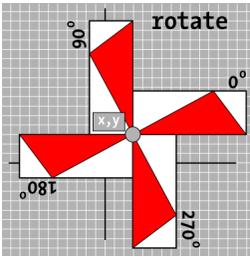
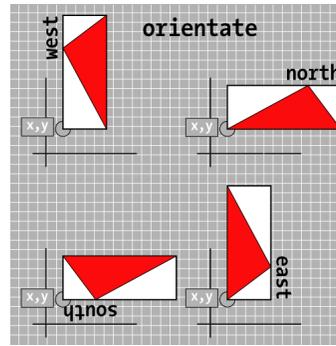


Fig. 8.4
The orientate option



With *fitmethod=meet* it is guaranteed that the image is not distorted and that it is placed in the box as large as possible.

Completely fitting the image into a box. We can further fit the image so that it completely fills the box. This is accomplished with *fitmethod=entire*. However, this combination will rarely be useful since the image may be distorted (see figure 8.2f).

Clipping an image when fitting it into a box. Using another fit method (*fitmethod=clip*) we can clip the object if it exceeds the target box. We decrease the box size to a width and height of 30 units and position the image in its original size at the center of the box (see figure 8.2g).

By positioning the image at the center of the box, the image will be cropped evenly on all sides. Similarly, to completely show the upper right part of the image you can position it with *position={right top}* (see figure 8.2h).

8.4.3 Orientating an Object

Placing an image with orientation. In our next example we orientate an image towards western direction (*orientate=west*). This means that the image is rotated by 90° counterclockwise and then the lower left corner of the rotated object is translated to the reference point (o, o) . The object will be rotated in itself (see figure 8.5a). Since we have not specified any fit method the image is output in its original size and exceeds the box.

Fitting an image proportionally into a box with orientation. Our next goal is to orientate the image to the west with a predefined size. We define a box of the desired size and fit the image into the box with the image's proportions being unchanged (*fitmethod=meet*). The orientation is specified as *orientate=west*. By default, the image will be placed in the lower left corner of the box (see figure 8.5b). figure 8.5c shows the image orientated to the east, and figure 8.5d the orientation to the south.

The *orientate* option supports the direction keywords *north*, *east*, *west*, and *south* as demonstrated in figure 8.4.

Note that the *orientate* option has no influence on the whole coordinate system but only on the placed object.

Fig. 8.5 Orientating an image

Generated output	Option list for <code>PDF_fit_image()</code>
a) 	<code>boxsize={70 45} orientate=west¹</code>
b) 	<code>boxsize={70 45} orientate=west fitmethod=meet</code>
c) 	<code>boxsize={70 45} orientate=east fitmethod=meet</code>
d) 	<code>boxsize={70 45} orientate=south fitmethod=meet</code>
e) 	<code>boxsize={70 45} position={center bottom} orientate=east fitmethod=clip</code>

1. The `boxsize` option isn't actually required because of the default `fitmethod=nofit`.

Fitting an oriented image into a box with clipping. We orientate the image to the east (`orientate=east`) and position it centered at the bottom of the box (`position={center bottom}`). In addition, we place the image in its original size and clip it if it exceeds the box (`fitmethod=clip`) (see figure 8.5e).

8.4.1 Rotating an Object

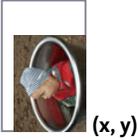
The `rotate` option rotates an object by rotating the coordinate system at the reference point. As a result the fitbox is also rotated. figure 8.3 demonstrates the general behavior of the `rotate` option.

Placing an image with rotation. Our first goal is to rotate an image by 90° counterclockwise. Before placing the object the coordinate system is rotated at the reference point by 90° counterclockwise. The rotated object's lower right corner (which is the unrotated object's lower left corner) ends up at the reference point. This case is shown in figure 8.6a.

Since the rotation affects the whole coordinate system, the box will be rotated as well. Similarly, we can rotate the image by 30° counterclockwise (see figure 8.6b).

Fitting an image with rotation. Our next goal is to fit the image rotated by 90° counterclockwise into the box while maintaining its proportions. This is accomplished using `fitmethod=meet` (see figure 8.6c). Similarly, we can rotate the image by 30° counterclockwise and proportionally fit the image into the box (see figure 8.6d).

Fig. 8.6 Rotating an image

Generated output	Option list for <code>PDF_fit_image()</code>
a) 	<code>boxsize={70 45} rotate=90¹</code>
b) 	<code>boxsize={70 45} rotate=30¹</code>
c) 	<code>boxsize={70 45} rotate=90 fitmethod=meet</code>
d) 	<code>boxsize={70 45} rotate=30 fitmethod=meet</code>

1. The `boxsize` option isn't actually required because of the default `fitmethod=nofit`.

8.4.2 Adjusting the Page Size

Adjusting the page size to an image. In the next example we will automatically adjust the page size to the object's size. This can be useful, for example, for archiving images in the PDF format. The reference point (x, y) can be used to specify whether the page will have exactly the object's size, or somewhat larger or smaller. When enlarging the page size (see figure 8.7) some border will be kept around the image. If the page size is smaller than the image some parts of the image will be clipped. Let's start with exactly matching the page size to the object's size:

```
p.fit_image(image, 0, 0, "adjustpage");
```

The next code fragment increases the page size by 40 units in x and y direction, creating a white border around the object:

```
p.fit_image(image, 40, 40, "adjustpage");
```

The next code fragment decreases the page size by 40 units in x and y direction. The object will be clipped at the page borders, and some area within the object (with a width of 40 units) will be invisible:

```
p.fit_image(image, -40, -40, "adjustpage");
```

In addition to placing by means of x and y coordinates (which specify the object's distance from the page edges, or the coordinate axes in the general case) you can also spec-

Fig. 8.7
Adjusting the page
size. Left to right:
exact, enlarge,
shrink



ify a target box. This is a rectangular area in which the object will be placed subject to various formatting rules. These can be controlled with the *boxsize*, *fitmethod* and *position* options.

Cloning the page boxes of an imported PDF page. You can copy all page boxes (MediaBox, CropBox) etc. of an imported PDF page to the current output page. The *cloneboxes* option must be supplied to *PDF_open_pdi_page()* to read all relevant box values, and again in *PDF_fit_pdi_page()* to apply the box values to the current page:

```
/* Open the page and clone the page box entries */
inpage = p.open_pdi_page(indoc, 1, "cloneboxes");
...
/* Start the output page with a dummy page size */
p.begin_page_ext(10, 10, "");
...
/*
 * Place the imported page on the output page, and clone all
 * page boxes which are present in the input page; this will
 * override the dummy size used in begin_page_ext().
 */
p.fit_pdi_page(inpage, 0, 0, "cloneboxes");
```

Using this technique you can make sure that the pages in the generated PDF will have the exact same page size, cropping etc. as the pages of the imported document. This is especially important for prepress applications.

8.4.1 Querying Information about placed Images and PDF Pages

Information about placed images and templates. The *PDF_info_image()* method can be used to query image and template information. The supported keywords for this method cover general image information (e.g. width and height in pixels) as well as geometry information related to placing the image on the output page (e.g. width and height in absolute values after performing the fitting calculations).

The following code fragment retrieves both the pixel size and the absolute size after placing an image with certain fitting options:

```
String optlist = "boxsize={300 400} fitmethod=meet orientate=west";
p.fit_image(image, 0.0, 0.0, optlist);

imagewidth = (int) p.info_image(image, "imagewidth", optlist);
imageheight = (int) p.info_image(image, "imageheight", optlist);
System.err.println("image size in pixels: " + imagewidth + " x " + imageheight);

width = p.info_image(image, "width", optlist);
height = p.info_image(image, "height", optlist);
System.err.println("image size in points: " + width + " x " + height);
```

Information about placed PDF pages. The `PDF_info_pdi_page()` method can be used to query information about placed PDF pages. The supported keywords for this method cover information about the original page (e.g. its width and height) as well as geometry information related to placing the imported PDF on the output page (e.g. width and height after performing the fitting calculations).

The following code fragment retrieves both the original size of the imported page and the size after placing the page with certain fitting options:

```
String optlist = "boxsize={400 500} fitmethod=meet";
p.fit_pdi_page(page, 0, 0, optlist);

pagewidth = p.info_pdi_page(page, "pagewidth", optlist);
pageheight = p.info_pdi_page(page, "pageheight", optlist);
System.err.println("original page size: " + pagewidth + " x " + pageheight);

width = p.info_pdi_page(page, "width", optlist);
height = p.info_pdi_page(page, "height", optlist);
System.err.println("size of placed page: " + width + " x " + height);
```

9 Text and Table Formatting

9.1 Placing and Fitting Textlines

The method `PDF_fit_textline()` for placing a single line of text on a page offers a wealth of formatting options. The most important options will be discussed in this section using some common application examples. A complete description of these options can be found in the PDFlib API Reference. Most options for `PDF_fit_textline()` are identical to those of `PDF_fit_image()`. Therefore we only use text-related examples here; it is recommended to take a look at the examples in Section 8.4, »Placing Images, Graphics, and imported PDF Pages«, page 190, for an introduction to image formatting.

The examples below demonstrate only the relevant call of `PDF_fit_textline()`, assuming that the required font has already been loaded and set in the desired font size.

`PDF_fit_textline()` uses a hypothetical text box to determine the positioning of the text: the width of the text box is identical to the width of the text, and the box height is identical to the height of capital letters in the font. The text box can be modified by the `matchbox` option.

In the examples below, the coordinates of the reference point are supplied as `x, y` parameters of `PDF_fit_textline()`. The fitbox for text lines is the area where text will be placed. It is defined as the rectangular area specified with the `x, y` parameters of `PDF_fit_textline()` and appropriate options (`boxsize, fitmethod, position, rotate`). The fitbox can be reduced to the left/right or top/bottom with the `margin` option.

Cookbook Code samples regarding text output issues can be found in the `text_output` category of the *PDFlib Cookbook*.

9.1.1 Simple Textline Placement

Positioning text at the reference point. By default, the text will be placed with the lower left corner at the reference point. However, in this example we want to place the text with the bottom centered at the reference point. The following code fragment places the text box with the bottom centered at the reference point (`30, 20`).

```
p.fit_textline(text, 30, 20, "position={center bottom}");
```

figure 9.1 illustrates centered text placement. Similarly, you can use the `position` option with another combination of the keywords `left, right, center, top, and bottom` to place text at the reference point.

Fig. 9.1
Centered text

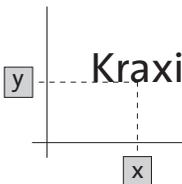
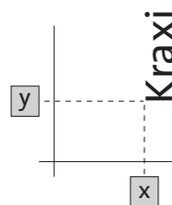


Fig. 9.2
Simple text with
orientation west



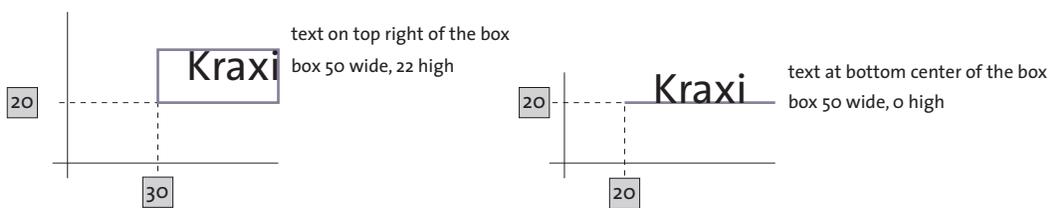


Fig. 9.3 Positioning text in a box

Placing text with orientation. Our next goal is to rotate text while placing its lower left corner (after the rotation) at the reference point. The following code fragment orientates the text to the west (90° counterclockwise) and then translates the lower left corner of the rotated text to the reference point $(0, 0)$.

```
p.fit_textline(text, 0, 0, "orientate=west");
```

figure 9.2 illustrates simple text placement with orientation.

9.1.1 Positioning Text in a Box

In order to position the text, an additional box with predefined width and height can be used, and the text can be positioned relative to this box. figure 9.3 illustrates the general behavior

Positioning text in the box. We define a rectangular box and place the text within this box on the top right. The following code fragment defines a box with a width of 50 units and a height of 22 units at reference point $(30, 20)$. In figure 9.4a, the text is placed on the top right of the box.

Similarly, we can place the text at the center of the bottom. This case is illustrated in figure 9.4b.

To achieve some distance between the text and the box we can add the *margin* option (see figure 9.4c).

Note that the box or line depicted for visualizing the box size in the figures is not part of the actual output.

Fig. 9.4 Placing text in a box subject to various positioning options

Generated output	Option list for <code>PDF_fit_textline()</code>
a) 	<code>boxsize={50 22} position={right top}</code>
b) 	<code>boxsize={50 22} position={center bottom}</code>
c) 	<code>boxsize={50 22} position={center bottom} margin={0 3}</code>
d) 	<code>boxsize={50 0} position={center bottom}</code>

Generated output	Option list for <code>PDF_fit_textline()</code>
e) 	<code>boxsize={0 35} position={left center} orientate=west</code>

Aligning text at a horizontal or vertical line. Positioning text along a horizontal or vertical line (i.e. a box with zero height or width) is a somewhat extreme case which may be useful nevertheless. In figure 9.4d the text is placed with the bottom centered at the box. With a width of 50 and a height of 0, the box resembles to a horizontal line.

To align the text centered along a vertical line we will orientate it to the west and position it at the left center of the box. This case is shown in figure 9.4e.

9.1.2 Fitting Text into a Box

In this section we use various fit methods to fit the text into the box. The current font and font size are assumed to be the same in all examples so that we can see how the font size and other properties will implicitly be changed by the different fit methods.

Let's start with the default case: no fit method will be used so that no clipping or scaling occurs. The text will be placed in the center of the box which is 100 units wide and 35 units high (see figure 9.5a).

Decreasing the box width from 100 to 50 units doesn't have any effect on the output. The text will remain in its original font size and will exceed beyond the box (see figure 9.5b).

Proportionally fitting text into a small box. Now we will completely fit the text into the box while maintaining its proportions. This can be achieved with the option `fitmethod=auto`. In figure 9.5c the box is wide enough to keep the text in its original size completely so that the text is placed into the box unchanged.

When scaling down the width of the box from 100 to 58, the text is too long to fit completely. The `auto` fit method tries to condense the text horizontally, subject to the `shrinklimit` option (default: 0.75). figure 9.5d shows the text shrunk to 75 percent of its original length.

When reducing the box width further to 30 units the text will not fit even with shrinking. Then the `meet` method is applied. It decreases the font size until the text fits completely into the box. This case is shown in figure 9.5e.

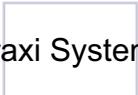
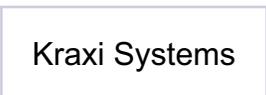
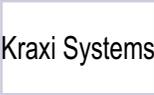
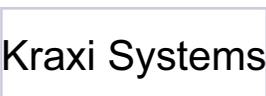
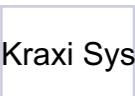
Fitting the text into the box with increased font size. You might want to fit the text so that it covers the whole width (or height) of the box but maintains its proportions. Using `fitmethod=meet` with a box larger than the text, the text will be increased until its width matches the box width. This case is illustrated in figure 9.5f.

Completely fitting text into a box. We can further fit the text so that it completely fills the box. In this case, `fitmethod=entire` is used. However, this combination will rarely be used since the text will most probably be distorted (see figure 9.5g).

Fitting text into a box with clipping. In another rare case you might want to fit the text in its original size and clip the text if it exceeds the box. In this case, `fitmethod=clip`

can be used. In figure 9.5h the text is placed at the bottom left of a box which is not broad enough. The text will be clipped on the right.

Fig. 9.5 Fitting text into a box on the page subject to various options

Generated output	Option list for <code>PDF_fit_textline()</code>
a) 	<code>boxsize={100 35} position=center fontsize=12</code>
b) 	<code>boxsize={50 35} position=center fontsize=12</code>
c) 	<code>boxsize={100 35} position=center fontsize=12 fitmethod=auto</code>
d) 	<code>boxsize={58 35} position=center fontsize=12 fitmethod=auto</code>
e) 	<code>boxsize={30 35} position=center fontsize=12 fitmethod=auto</code>
f) 	<code>boxsize={100 35} position=center fontsize=12 fitmethod=meet</code>
g) 	<code>boxsize={100 35} position=center fontsize=12 fitmethod=entire</code>
h) 	<code>boxsize={50 35} position={left center} fontsize=12 fitmethod=clip</code>

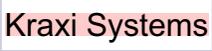
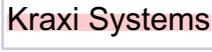
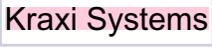
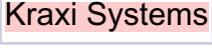
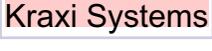
Vertically centering text. The text height in `PDF_fit_textline()` is the capheight, i.e. the height of the capital letter *H*, by default. If the text is positioned in the center of a box it will be vertically centered according to its capheight (see figure 9.6a).

To specify another height for the text box we can use the Matchbox feature (see also Section 9.4, »Matchboxes«, page 241). The `matchbox` option of `PDF_fit_textline()` define the height of a Textline which is the capheight of the given font size, by default. The height of the matchbox is calculated according to its `boxheight` suboption. The `boxheight` suboption determines the extent of the text above and below the baseline. `matchbox={boxheight={capheight none}}` is the default setting, i.e. the top border of the matchbox will touch the capheight above the baseline, and the bottom border of the matchbox will not extend below the baseline.

To illustrate the size of the matchbox we will fill it with red color (see figure 9.6b). figure 9.6c vertically centers the text according to the *xheight* by defining a matchbox with a corresponding box height.

figure 9.6d–f shows the matchbox (red) with various useful *boxheight* settings to determine the height of the text to be centered in the box.

Fig. 9.6 Fitting text proportionally into a box according to different box heights

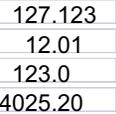
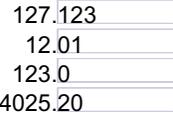
Generated output	Option list for PDF_fit_textline()
a) 	<code>boxsize={80 20} position=center fitmethod=auto</code>
b) 	<code>boxsize={80 20} position=center fitmethod=auto matchbox={boxheight={capheight none} fillcolor=mistyrose}</code>
c) 	<code>boxsize={80 20} position=center fitmethod=auto matchbox={boxheight={xheight none} fillcolor=mistyrose}</code>
d) 	<code>boxsize={80 20} position=center fitmethod=auto matchbox={boxheight={ascender none} fillcolor=mistyrose}</code>
e) 	<code>boxsize={80 20} position=center fitmethod=auto matchbox={boxheight={ascender descender} fillcolor=mistyrose}</code>
f) 	<code>boxsize={80 20} position=center fitmethod=auto matchbox={boxheight={fontsize none} fillcolor=mistyrose}</code>

9.1.3 Aligning Text at a Character

You might want to align text at a certain character, e.g. at the decimal point in a number. As shown in figure 9.7a, the text is positioned at the center of the fitbox. Using `PDF_fit_textline()` with the `alignchar=.` option the numbers are aligned at the dot character.

You can omit the `position` option which places the dots in the center of the box. In this case, the default `position={left bottom}` will be used which places the dots at the reference point (see figure 9.7b). In general, the alignment character will be placed with the lower right corner at the reference point.

Fig. 9.7 Aligning a Textline to the dot character

Generated output	Option list for PDF_fit_textline()
a) 	<code>boxsize={70 8} position={center bottom} alignchar=.</code>
b) 	<code>boxsize={70 8} position={left bottom} alignchar=.</code>

9.1.4 Placing a Stamp

Cookbook A full code sample can be found in the *Cookbook* topic `text_output/simple_stamp`.

As an alternative to rotated text, the stamp feature offers a convenient method for placing text diagonally in a box. The stamp method will automatically perform some sophisticated calculations to determine a suitable font size and rotation so that the text covers the box. To place a diagonal stamp, e.g. in the page background, use `PDF_fit_textline()` with the `stamp` option. With `stamp=ll2ur` the text is placed from the lower left to the upper right corner of the fitbox. However, with `stamp=ul2lr` the text is placed from the upper left to the lower right corner of the fitbox. The `fitbox` option is ignored. As shown in figure 9.8, `showborder=true` is used to illustrate the fitbox and the bounding box of the stamp.

Fig. 9.8 Fitting a text line like a stamp from the lower left to the upper right

Generated output	Option list for <code>PDF_fit_textline()</code>
	<code>fontsize=8 boxsize={160 50} stamp=ll2ur showborder=true</code>

9.1.5 Using Leaders

Leaders can be used to fill the space between the borders of the fitbox and the text. For example, dot leaders are often used as a visual aid between the entries in a table of contents and the corresponding page numbers.

Leaders in a table of contents. Using `PDF_fit_textline()` with the `leader` option and the `alignment={none right}` suboption, leaders are appended to the right of the text line, and repeated until the right border of the text box. There will be an equal distance between the rightmost leader and the right border, while the distance between the text and the leftmost leader may differ (see figure 9.9a).

Cookbook A full code sample demonstrating the usage of dot leaders in a text line can be found in the *Cookbook* topic `textoutput/leaders_in_textline`.

Cookbook A full code sample demonstrating the usage of dot leaders in a Textflow can be found in the *Cookbook* topic `textflow/dot_leaders_with_tabs`.

Leaders in a news ticker. In another use case you might want to create a news ticker effect. In this case we use a plus and a space character `»+ «` as leaders. The text line is placed in the center, and the leaders are printed before and after the text line (`alignment={left right}`). The left and right leaders are aligned to the left and right border, and might have a varying distance to the text (see figure 9.9b).

Fig. 9.9 Fitting a text line using leaders

Generated output	Option list for <code>PDF_fit_textline()</code>
<p>a)</p> <pre> Features of Giant Wing Description of Long Distance Glider..... Benefits of Cone Head Rocket</pre>	<pre> boxsize={200 10} leader={alignment={none right}}</pre>
<p>b)</p> <pre> +++++ Giant Wing in purple!+++++ ++ Long Distance Glider with sensational range! ++ +++++ Cone Head Rocket incredibly fast! +++++</pre>	<pre> boxsize={200 10} position={center bottom} leader={alignment={left right}} text={+ }</pre>

9.1.6 Text on a Path

Instead of placing text on a straight line you can also place text on an arbitrary path. PDFlib will place the individual characters along the path so that the text follows the curvature of the path. Use the `textpath` option of `PDF_fit_textline()` to create text on a path. The path must have been created earlier and is represented by a path handle. Path handles can be created by explicitly constructing a path with `PDF_add_path_point()` and related path object functions, or by retrieving a handle for the clipping path in an existing raster image. The following code fragment creates a simple path and places text on the path (see figure 9.10):

```

/* Define the path in the origin */
path = p.add_path_point( -1, 0, 0, "move", "");
path = p.add_path_point(path, 100, 100, "control", "");
path = p.add_path_point(path, 200, 0, "circular", "");

/* Place text on the path */
p.fit_textline("Long Distance Glider with sensational range!", x, y,
    "textpath={path=" + path + "} position={center bottom}");

/* We also draw the path for demonstration purposes */
p.draw_path(path, x, y, "stroke strokecolor=dodgerblue");

```

Cookbook A full code sample can be found in the Cookbook topic `text_output/text_on_a_path`.

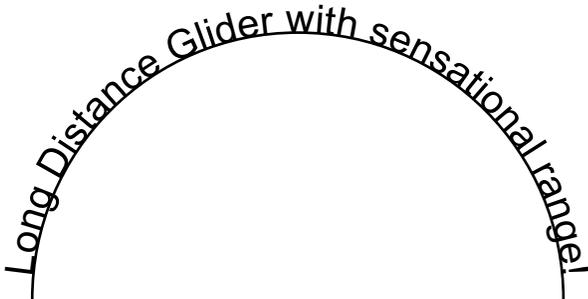


Fig. 9.10
Text on a path

Using an image clipping path for placing text. As an alternative to manually constructing a path object with the path functions you can extract the clipping path from an image and place text on the resulting path. The image must have been loaded with the *honorclippingpath* option, and the *clippingpathname* option must also be supplied to *PDF_load_image()* if the target path is not the image's default clipping path:

```
image = p.load_image("auto", "image.tif", "clippingpathname={path 1}");

/* create a path object from the image's clipping path */
path = (int) p.info_image(image, "clippingpath", "");
if (path == -1)
    throw new Exception("Error: clipping path not found!");

/* Place text on the path */
p.fit_textline("Long Distance Glider with sensational range!", x, y,
    "textpath={path=" + path + "} position={center bottom}");
```

Creating a gap between path and text. By default, PDFlib places individual characters directly on the path, i.e. there is no space between the glyphs and the path. If you want to create a gap between the path and the text you can increase the character boxes. This can be achieved with *boxheight* suboption of the *matchbox* option which specifies the vertical extension of the character boxes. The following option list takes the descenders into account (see figure 9.11):

```
p.fit_textline("Long Distance Glider with sensational range!", x, y,
    "textpath={path=" + path + "} position={center bottom} " +
    "matchbox={boxheight={capheight descender}}");
```

9.1.7 Shadowed Text

The *shadow* option can be used to create a shadow effect for text. You can specify the shadow color as well as its horizontal and vertical distance from the main text in suboptions:

```
p.fit_textline("Long Distance Glider", x, y,
    "fillcolor=rosybrown shadow={offset={3, -3}}");
```

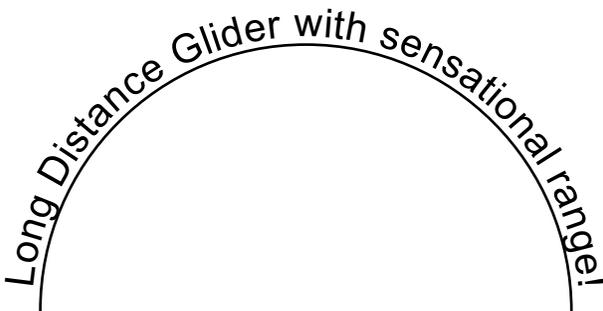
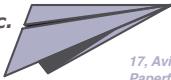


Fig. 9.11
Text on a path with an additional
gap between text and path

9.2 Multi-Line Textflows

In addition to placing single lines of text on the page, PDFlib supports a powerful feature called Textflow which can be used to format arbitrarily long text portions. The text may extend across any number of lines, columns, or pages, and its appearance can be controlled with a variety of options. Character properties such as font, size, and color can be applied to any part of the text. Textflow properties such as justified or ragged text, paragraph indentation and tab stops can be specified; line breaking opportunities designated by soft hyphens in the text will be taken into account. You can also supply tags with structure information. figure 9.12 and figure 9.13 demonstrate how various parts of an invoice can be placed on the page using the Textflow feature. We will discuss the options for controlling the output in more detail in the following sections.

Kraxi Systems, Inc.
Paper Planes



17, Aviation Road
Paperfield
Phone 7079-4301
Fax 7079-4302
info@kraxi.com
www.kraxi.com

Kraxi Systems, Inc. 17, Aviation Road Paperfield
John Q. Doe
255 Customer Lane
Suite B
12345 User Town
Everland

INVOICE 14.03.2004

ITEM	DESCRIPTION	QUANTITY	PRICE	AMOUNT
1	Super Kite	2	20,00	40,00
2	Turbo Flyer	5	40,00	200,00
3	Giga Trash	1	180,00	180,00
4	Bare Bone Kit	3	50,00	150,00
5	Nitty Gritty	10	20,00	200,00
6	Pretty Dark Flyer	1	75,00	75,00
7	Free Gift	1	0,00	0,00
				845,00

Terms of payment: 30 days net. 30 days warranty starting at the day of sale. This warranty covers defects in workmanship only. Kraxi Systems, Inc., at its option, repairs or replaces the product under warranty. This warranty is not transferable. Returns or exchanges are not possible for wet products.

Have a look at our new paper plane models!
 Our paper planes are the ideal way of passing the time. We offer revolutionary new developments of the traditional common paper planes. If your lesson, conference, or lecture turn out to be deadly boring, you can have a wonderful time with our planes. All our models are folded from one paper sheet.
 They are exclusively folded without using any adhesive. Several models are equipped with a folded landing gear enabling a safe landing on the intended location provided that you have aimed well. Other models are able to fly loops or cover long distances. Let them start from a vista point in the mountains and see where they touch the ground.

1. Long Distance Glider
 With this paper rocket you can send all your messages even when sitting in a hall or in the cinema pretty near the back.

2. Giant Wing
 An unbelievable sailplane! It is amazingly robust and can even do

Fig. 9.12
Formatting
Textflows

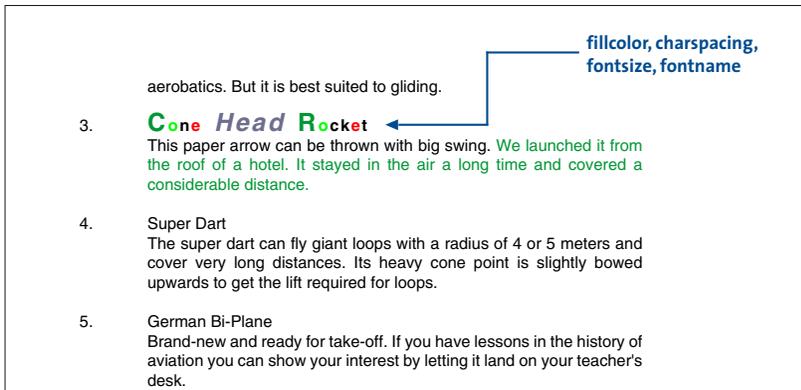


Fig. 9.13
Formatting
Textflows

A multi-line Textflow can be placed into one or more rectangles (so-called fitboxes) on one or more pages. The following steps are required for placing a Textflow on the page:

- ▶ The method `PDF_add_textflow()` accepts portions of text and corresponding formatting options (where the options can be provided as a separate string or inline with the text), creates a Textflow object, and returns a handle. As an alternative, the method `PDF_create_textflow()` analyzes the complete text in a single call, where the text may contain inline options for formatting control. These functions do not place any text on the page.
- ▶ The method `PDF_fit_textflow()` places all or parts of the Textflow in the supplied fitbox. To completely place the text, this step must be repeated where each of the method calls provides a new fitbox which may be located on the same or another page.
- ▶ The method `PDF_delete_textflow()` deletes the Textflow object after it has been placed in the document.

The functions `PDF_add/create_textflow()` for creating Textflows support a variety of options for controlling the formatting process. These options can be provided in the method's option list, or embedded as inline options in the text. `PDF_info_textflow()` can be used to query formatting results and many other Textflow details. We will discuss Textflow placement using some common application examples. A complete list of Textflow options can be found in the *PDFlib API Reference*.

Many of the options supported by `PDF_add/create_textflow()` are identical to those of `PDF_fit_textline()`. It is therefore recommended to familiarize yourself with the examples in Section 9.1, »Placing and Fitting Textlines«, page 197. In the below sections we will focus on options related to multi-line text.

Cookbook Code samples regarding text output issues can be found in the `text_output` category of the *PDFlib Cookbook*.

9.2.1 Placing Textflows in the Fitbox

The fitbox for Textflow is the area where text will be placed. It is defined as the rectangular area specified with the `llx`, `lly`, `urx`, `ury` parameters of `PDF_fit_textflow()`.

Placing text in a single fitbox. Let's start with an easy example. The following code fragment uses two calls to `PDF_add_textflow()` to assemble a piece of bold text and a piece of normal text. Font and font size are specified explicitly. In the first call to `PDF_add_textflow()`, `-1` is supplied, and the Textflow handle will be returned to be used in subsequent calls to `PDF_add_textflow()`, if required. `text1` and `text2` are assumed to contain the actual text to be printed.

With `PDF_fit_textflow()`, the resulting Textflow is placed in a fitbox on the page using default formatting options.

```
/* Add text with bold font */
tf = p.add_textflow(-1, text1, "fontname=Helvetica-Bold fontsize=9");
if (tf == -1)
    throw new Exception("Error: " + p.get_errmsg());

/* Add text with normal font */
tf = p.add_textflow(tf, text2, "fontname=Helvetica fontsize=9");
if (tf == -1)
    throw new Exception("Error: " + p.get_errmsg());

/* Place all text */
result = p.fit_textflow(tf, left_x, left_y, right_x, right_y, "");
if (!result.equals("_stop"))
    { /* ... */ }

p.delete_textflow(tf);
```

Placing text in two fitboxes on multiple pages. If the text placed with `PDF_fit_textflow()` doesn't completely fit into the fitbox, the output will be interrupted and the method will return the string `_boxfull`. PDFlib will remember the amount of text already placed, and will continue with the remainder of the text when the method is called again. In addition, it may be necessary to create a new page. The following code fragment demonstrates how to place a Textflow in two fitboxes per page on one or more pages until the text has been placed completely (see figure 9.14).

Cookbook A full code sample can be found in the Cookbook topic `textflow/starter_textflow`.



Fig. 9.14
Placing a Textflow in fitboxes

```

/* Loop until all of the text is placed; create new pages as long as more text needs
 * to be placed. Two columns will be created on all pages.
 */
do
{
    String optlist = "verticalalign=justify linespreadlimit=120%";

    p.begin_page_ext(0, 0, "width=a4.width height=a4.height");

    /* Fill the first column */
    result = p.fit_textflow(tf, llx1, lly1, urx1, ury1, optlist);

    /* Fill the second column if we have more text*/
    if (!result.equals("_stop"))
        result = p.fit_textflow(tf, llx2, lly2, urx2, ury2, optlist);

    p.end_page_ext("");

    /* "_boxfull" means we must continue because there is more text;
     * "_nextpage" is interpreted as "start new column"
     */
} while (result.equals("_boxfull") || result.equals("_nextpage"));

/* Check for errors */
if (!result.equals("_stop"))
{
    /* "_boxempty" happens if the box is very small and doesn't hold any text at all.
     */
    if (result.equals( "_boxempty"))
        throw new Exception("Error: " + p.get_errmsg());
    else
    {
        /* Any other return value is a user exit caused by the "return" option;
         * this requires dedicated code to deal with.
         */
    }
}
p.delete_textflow(tf);

```

9.2.2 Paragraph Formatting Options

In the previous example we used default settings for the paragraphs. For example, the default alignment is left-justified, and the leading is 100% (which equals the font size).

In order to fine-tune the paragraph formatting we can feed more options to `PDF_add_textflow()`. For example, we can indent the text 15 units from the left and 10 units from the right margin. The first line of each paragraph should be indented by an additional 20 units. The text should be justified against both margins, and the leading increased to 140%. Finally, we'll reduce the font size to 8 points. To achieve this, extend the option list for `PDF_add_textflow()` as follows (see figure 9.15):

```

String optlist =
    "leftindent=15 rightindent=10 parindent=20 alignment=justify " +
    "leading=140% fontname=Helvetica fontsize=8";

```

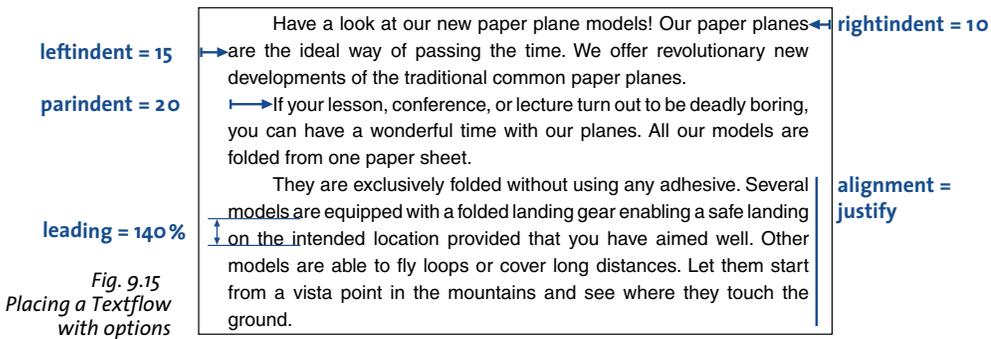


Fig. 9.15
Placing a Textflow
with options

9.2.1 Inline Option Lists and Macros

The text in figure 9.15 is not yet perfect. The headline »Have a look at our new paper plane models!« should sit on a line of its own, should use a larger font, and should be centered. There are several ways to achieve this.

Inline option lists. Up to now we provided formatting options in an option list supplied directly to the method. In order to continue the same way we would have to split the text, and place it in two separate calls, one for the headline and another one for the remaining text. However, in certain situations, e.g. with lots of formatting changes, this method may be cumbersome.

To facilitate option handling you can supply inline options for a Textflow, i.e. text contents and formatting options are mixed as required. `PDF_create_textflow()` and `PDF_add_textflow()` with the option `inlineoptions` interpret text and inline options which are embedded directly in the text. Inline option lists are provided as part of the body text. By default, they are delimited by »<<« and »>>« characters. We will therefore integrate the options for formatting the heading and the remaining paragraphs into our body text as follows.

Note Inline option lists are colored in all subsequent samples; end-of-paragraph characters are visualized with arrows.

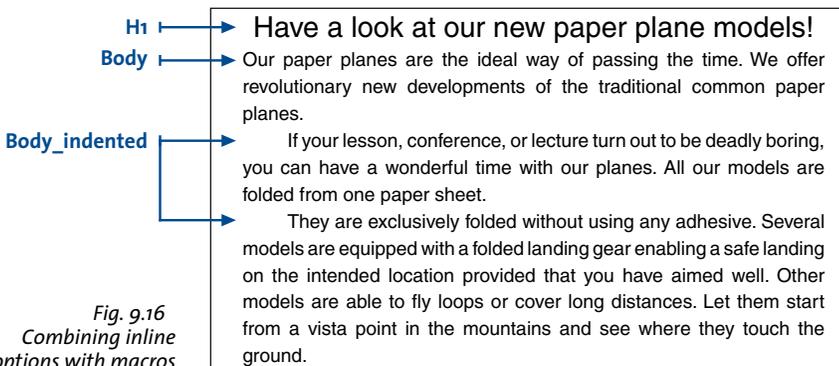


Fig. 9.16
Combining inline
options with macros

```
<leftindent=15 rightindent=10 alignment=center fontname=Helvetica fontsize=12>
Have a look at our new paper plane models! ←
<alignment=justify fontname=Helvetica leading=140% fontsize=8>
Our paper planes are the ideal way of passing the time. We offer
revolutionary new developments of the traditional common paper planes. ←
<parindent=20>If your lesson, conference, or lecture
turn out to be deadly boring, you can have a wonderful time
with our planes. All our models are folded from one paper sheet. ←
They are exclusively folded without using any adhesive. Several
models are equipped with a folded landing gear enabling a safe
landing on the intended location provided that you have aimed well.
Other models are able to fly loops or cover long distances. Let them
start from a vista point in the mountains and see
where they touch the ground.
```

By default, `PDF_add_textflow()` searches for inline option lists, where the characters for bracketing option lists can be redefined with the `begoptlistchar` and `endoptlistchar` options. Supplying the keyword `none` for the `begoptlistchar` option completely disables the search for option lists. This is useful if the text doesn't contain any inline option lists, and you want to make sure that »« and »»« will be processed as regular characters.

`PDF_add_textflow()` searches for inline option lists only if the option `inlineoptions` is provided.

Symbol characters and inline option lists. Symbolic characters can be used for Textflow even in combination with inline option lists. The code for the character which introduces an inline option list (by default: 'U+003C) will not be interpreted as a symbol code within text for a font with `encoding=builtin`. In order to select the symbol glyph with the same code, the same workarounds can be used which are available for text fonts, i.e. by redefining the start character with the `begoptlistchar` option or by using the `textlen` option to specify the number of symbolic glyphs. Note that character references (e.g. `<`) can not be used as a workaround. Also note that the `textlen` option is not available for `PDF_add_textflow()` with the option `inlineoptions`.

Macros. The text above contains several different types of paragraphs, such as heading or body text with or without indentation. Each of these paragraph types is formatted differently and occurs multiply in longer Textflows. In order to avoid starting each paragraph with the corresponding inline options, we can combine these in macros, and refer to the macros in the text via their names. As shown in figure 9.16 we define three macros called `H1` for the heading, `Body` for main paragraphs, and `Body_indented` for indented paragraphs. In order to use a macro we place the `&` character in front of its name and put it into an option list. The following code fragment defines three macros according to the previously used inline options and uses them in the text:

```
<macro {
H1 {leftindent=15 rightindent=10 alignment=center
fontname=Helvetica fontsize=12}

Body {leftindent=15 rightindent=10 alignment=justify leading=140%
fontname=Helvetica fontsize=8}

Body_indented {parindent=20 leftindent=15 rightindent=10 alignment=justify
leading=140% fontname=Helvetica fontsize=8}
}>
```

```

<&H1>Have a look at our new paper plane models! ←
<&Body>Our paper planes are the ideal way of passing the time. We offer
revolutionary new developments of the traditional common paper planes. ←
<&Body_indented>If your lesson, conference, or lecture
turn out to be deaded boring, you can have a wonderful time
with our planes. All our models are folded from one paper sheet. ←
They are exclusively folded without using any adhesive. Several
models are equipped with a folded landing gear enabling a safe
landing on the intended location provided that you have aimed well.
Other models are able to fly loops or cover long distances. Let them
start from a vista point in the mountains and see
where they touch the ground.

```

Explicitly setting options. Note that all options which are not set in macros will retain their previous values. In order to avoid side effects caused by unwanted »inheritance« of options you should explicitly specify all settings required for a particular macro. This way you can ensure that the macros will behave consistently regardless of their ordering or combination with other option lists.

On the other hand, you can take advantage of this behavior for deliberately retaining certain settings from the context instead of supplying them explicitly. For example, a macro could specify the font name without supplying the *fontsize* option. As a result, the font size will always match that of the preceding text.

Inline options or options passed as method parameters? When using Textflows it makes an important difference whether the text is contained literally in the program code or comes from some external source, and whether the formatting instructions are separate from the text or part of it. In most applications the actual text will come from some external source such as a database. In practice there are two main scenarios:

- ▶ Text contents from external source, formatting options in the program: An external source delivers small text fragments which are assembled within the program, and combined with formatting options (in the method call) at runtime.
- ▶ Text contents and formatting options from external source: Large amounts of text including formatting options come from an external source. The formatting is provided by inline options in the text, represented as simple options or macros. When it comes to macros a distinction must be made between macro definition and macro call. This allows an interesting intermediate form: the text content comes from an external source and contains macro calls for formatting. However, the macro definitions are only blended in at runtime. This has the advantage that the formatting can easily be changed without having to modify the external text. For example, when generating greeting cards one could define different styles via macros to give the card a romantic, technical, or other touch.

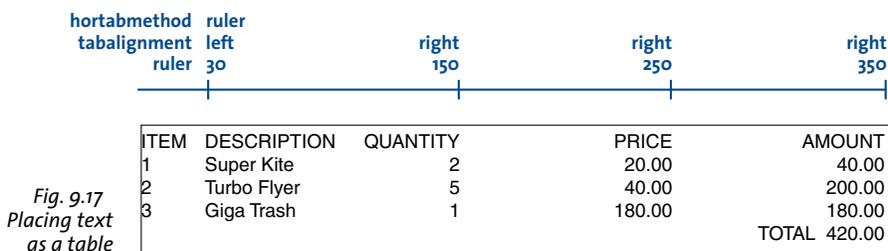
9.2.1 Tab Stops

In the next example we will place a simple table with left- and right-aligned columns using tab characters. The table contains the following lines of text, where individual entries are separated from each other with a tab character (indicated by arrows):

```

ITEM → DESCRIPTION → QUANTITY → PRICE → AMOUNT ←
1 → Super Kite → 2 → 20.00 → 40.00 ←
2 → Turbo Flyer → 5 → 40.00 → 200.00 ←

```



```
3 → Giga Trash → 1 → 180.00 → 180.00 ←
→ → → → TOTAL 420.00
```

To place that simple table use the following option list in `PDF_add/create_textflow()`. The `ruler` option defines the tab positions, `tabalignment` specifies the alignment of tab stops, and `hortabmethod` specifies the method used to process tab stops (the result can be seen in figure 9.17):

```
String optlist =
    "ruler      = {30    150  250  350} " +
    "tabalignment={left right right right} " +
    "hortabmethod=ruler leading=120% fontname=Helvetica fontsize=9";
```

Cookbook A full code sample can be found in the *Cookbook* topic `textflow/tabstops_in_text`.

Note PDFlib's table feature is recommended for creating complex tables (see Section 9.3, «Table Formatting», page 225).

9.2.1 Numbered Lists and Paragraph Spacing

The following example demonstrates how to format a numbered list using the inline option `leftindent` (see figure 9.18):

```
1.<leftindent 10>Long Distance Glider: With this paper rocket you can send all
your messages even when sitting in a hall or in the cinema pretty near the back. ←
<leftindent 0>2.<leftindent 10>Giant Wing: An unbelievable sailplane! It is amazingly
robust and can even do aerobatics. But it is best suited to gliding. ←
<leftindent 0>3.<leftindent 10>Cone Head Rocket: This paper arrow can be thrown with big
swing. We launched it from the roof of a hotel. It stayed in the air a long time and
covered a considerable distance.
```

Cookbook Full code samples for bulleted and numbered lists can be found in the *Cookbook* topics `textflow/bulleted_list` and `textflow/numbered_list`.

Setting and resetting the indentation value is cumbersome, especially since it is required for each paragraph. A more elegant solution defines a macro called `list`. For convenience it defines a macro `indent` which is used as a constant. The macro definitions are as follows:

```
<macro {
indent {25}

list {parindent=-&indent leftindent=&indent hortabsize=&indent
hortabmethod=ruler ruler={&indent}}
}>
```

`leftindent = &indent`
`parindent = -&indent`



1. Long Distance Glider: With this paper rocket you can send all your messages even when sitting in a hall or in the cinema pretty near the back.
2. Giant Wing: An unbelievable sailplane! It is amazingly robust and can even do aerobatics. But it is best suited to gliding.
3. Cone Head Rocket: This paper arrow can be thrown with big swing. We launched it from the roof of a hotel. It stayed in the air a long time and covered a considerable distance.

Fig. 9.19
Numbered list
with macros

```
<&list>1. → Long Distance Glider: With this paper rocket you can send all your messages  
even when sitting in a hall or in the cinema pretty near the back. ←  
2. → Giant Wing: An unbelievable sailplane! It is amazingly robust and can even do  
aerobatics. But it is best suited to gliding. ←  
3. → Cone Head Rocket: This paper arrow can be thrown with big swing. We launched  
it from the roof of a hotel. It stayed in the air a long time and covered a  
considerable distance.
```

The *leftindent* option specifies the distance from the left margin. The *parindent* option, which is set to the negative of *leftindent*, cancels the indentation for the first line of each paragraph. The options *hortabsize*, *hortabmethod*, and *ruler* specify a tab stop which corresponds to *leftindent*. It makes the text after the number to be indented with the amount specified in *leftindent*. figure 9.19 shows the *parindent* and *leftindent* options at work.

Setting the distance between two paragraphs. In many cases more distance between adjacent paragraphs is desired than between the lines within a paragraph. This can be achieved by inserting an extra empty line (which can be created with the *nextline* option), and specifying a suitable leading value for this empty line. This value is the distance between the baseline of the last line of the previous paragraph and the baseline of the empty line. The following example will create 80% additional space between the two paragraphs (where 100% equals the most recently set value of the font size):

```
1. → Long Distance Glider: With this paper rocket you can send all your messages  
even when sitting in a hall or in the cinema pretty near the back.  
<nextline leading=80%><nextparagraph leading=100%>2. → Giant Wing: An unbelievable  
sailplane! It is amazingly robust and can even do aerobatics. But it is best suited to  
gliding.
```

Cookbook A full code sample can be found in the *Cookbook topic* `textflow/`
`distance_between_paragraphs`.

1. Long Distance Glider: With this paper rocket you can send all your messages even when sitting in a hall or in the cinema pretty near the back.
2. Giant Wing: An unbelievable sailplane! It is amazingly robust and can even do aerobatics. But it is best suited to gliding.
3. Cone Head Rocket: This paper arrow can be thrown with big swing. We launched it from the roof of a hotel. It stayed in the air a long time and covered a considerable distance.

Fig. 9.18
Numbered list

9.2.1 Control Characters and Character Mapping

Control characters in Textflows. Various characters are given special treatment in Textflows. PDFlib supports symbolic character names which can be used instead of the corresponding character codes in the *charmapping* option (which replaces characters in the text before processing it, see below). Table 9.1 lists all control characters which are evaluated by the Textflow functions along with their symbolic names, and explains their meaning. An option must only be used once per option list, but multiple option lists can be provided one after the other. For example, the following sequence will create an empty line:

```
<newline><newline>
```

Table 9.1 Control characters and their meaning in Textflows

Unicode character	entity name	equiv. Text-flow option	meaning within Textflows
U+0020	SP, space	space	align words and break lines
U+00A0	NBSP, nbsp	(none)	(no-break space) space character which will not break lines
U+202F	NNBSP, nnbsp	(none)	(narrow no-break space) fixed-width space character which will not break lines, and will not change its width according to formatting options
U+0009	HT, hortab	(none)	horizontal tab: will be processed according to the ruler, tabalignchar, and tabalignment options
U+002D	HY, hyphen	(none)	separator character for hyphenated words
U+00AD	SHY, shy	(none)	(soft hyphen) hyphenation opportunity, only visible at line breaks
U+000B U+2028	VT, verttab LS, linesep	newline	(next line) forces a new line
U+000A U+000D U+000D and U+000A	LF, linefeed CR, return CRLF	next-paragraph	(next paragraph) Same effect as newline; in addition, the parindent option will affect the next line.
U+0085 U+2029	NEL, newline PS, parasep		
U+000C	FF, formfeed	return	PDF_fit_textflow() stops and returns the string _nextpage.

Mapping/removing characters or sequences of characters. The *charmapping* option can be used to map or remove some characters in the text to others. Let's start with an easy case where we will map all tabs in the text to space characters. The *charmapping* option to achieve this looks as follows:

```
charmapping={hortab space}
```

This command uses the symbolic character names *hortab* and *space*. To achieve multiple mappings at once you can use the following command which will replace all tabs and line break combinations with space characters:

```
charmapping={hortab space CRLF space LF space CR space}
```

The following command removes all soft hyphens:

```
charmapping={shy {shy 0}}
```

Each tab character will be replaced with four space characters:

```
charmapping={horthab {space 4}}
```

Each arbitrary long sequence of linefeed characters will be reduced to a single linefeed character:

```
charmapping={linefeed {linefeed -1}}
```

Each sequence of CRLF combinations will be replaced with a single space:

```
charmapping={CRLF {space -1}}
```

We will take a closer look at the last example. Let's assume you receive text where the lines have been separated with fixed line breaks by some other software, and therefore cannot be properly formatted. You want to replace the linebreaks with space characters in order to achieve proper formatting within the fitbox. To achieve this we replace arbitrarily long sequences of linebreaks with a single space character. The initial text looks as follows:

```
To fold the famous rocket looper proceed as follows: ← ←
Take a sheet of paper. Fold it ←
lengthwise in the middle. ←
Then, fold down the upper corners. Fold the ←
long sides inwards ←
that the points A and B meet on the central fold.
```

The following code fragment demonstrates how to replace the redundant linebreak characters and format the resulting text:

```
/* assemble option list */
String optlist =
    "fontname=Helvetica fontsize=9 alignment=justify " +
    "charmapping {CRLF {space -1}}"

/* place textflow in fitbox */
textflow = p.add_textflow(-1, text, optlist);
if (textflow == -1)
    throw new Exception("Error: " + p.get_errmsg());

result = p.fit_textflow(textflow, left_x, left_y, right_x, right_y, "");
if (!result.equals("_stop"))
    { /* ... */ }

p.delete_textflow(textflow);
```

figure 9.20 shows Textflow output with the unmodified text and the improved version with the *charmapping* option.

To fold the famous rocket looper proceed as follows:

Take a sheet of paper. Fold it lengthwise in the middle. Then, fold down the upper corners. Fold the long sides inwards that the points A and B meet on the central fold.

Fig. 9.20
Top: text with redundant line breaks

To fold the famous rocket looper proceed as follows: Take a sheet of paper. Fold it lengthwise in the middle. Then, fold down the upper corners. Fold the long sides inwards that the points A and B meet on the central fold.

Bottom: replacing the linebreaks with the charmapping option

9.2.2 Hyphenation

PDFlib does not automatically hyphenate text, but can break words at hyphenation opportunities which are explicitly marked in the text by soft hyphen characters. The soft hyphen character is at position `U+00AD` in Unicode. Alternatively a character entity reference can be used: `­`

The glyph `U+00AD` will be used as hyphenation character if it is available in the font, otherwise `U+002D`. In addition to breaking opportunities designated by soft hyphens, words can be forcefully hyphenated in extreme cases when other methods of adjustment, such as changing the word spacing or shrinking text, are not possible.

Justified text with or without hyphen characters. In the following example we will print the following text with justified alignment. The text contains soft hyphen characters (visualized here as dashes):

Our paper planes are the ideal way of pas - sing the time. We offer revolu - tionary brand new dev - elop - ments of the tradi - tional common paper planes. If your lesson, confe - rence, or lecture turn out to be deadly boring, you can have a wonder - ful time with our planes. All our models are folded from one paper sheet. They are exclu - sively folded without using any adhe - sive. Several models are equip - ped with a folded landing gear enab - ling a safe landing on the intended loca - tion provided that you have aimed well. Other models are able to fly loops or cover long dist - ances. Let them start from a vista point in the mount - ains and see where they touch the ground.

figure 9.21 shows the generated text output with default settings for justified text. It looks perfect since the conditions are optimal: the fitbox is wide enough, and there are explicit break opportunities specified by the soft hyphen characters. As you can see in figure 9.22, the output looks okay even without explicit soft hyphens. The option list in both cases looks as follows:

```
fontname=Helvetica fontsize=9 alignment=justify
```

9.2.3 Widow and Orphan Lines

If the first line (or lines) of a paragraph appears by itself at the bottom of a column or page it is called an orphan. Similarly, if the last line (or lines) of a paragraph appears at the beginning of the next column or page it is called a widow. Isolated orphan or widow lines are considered as undesirable in high-quality typesetting.

Orphan control. The Textflow option *minlinecount* specifies the minimum number of lines in the last paragraph of the fitbox. If there are fewer lines they are placed in the next fitbox. The value *minlinecount=2* can be used to prevent single orphan lines of a paragraph at the end of a fitbox.

Widow control. Since PDFlib doesn't know anything about future Textflow placements it cannot directly control widow lines. However, you can implement widow control in the client code according to the following scheme:

- ▶ Fit the first part of the Textflow into the first fitbox in blind mode (option *blind=true*), i.e. without creating any real output.
- ▶ Fit the next part of the Textflow into the second fitbox in blind mode. Call *PDF.info_textflow()* with the keyword *firstparalinecount* to query the number of lines in the first paragraph of the second fitbox. If the result is 1 you found a single-line widow.
- ▶ Now you can return to the first fitbox with the *rewind* option, and your algorithm must adjust the fitting options to avoid the widow line. For example, this can be achieved by reducing the number of lines in the first fitbox with the *maxlines* option.

Cookbook A full code sample can be found in the *Cookbook* topic *textflow/widows_and_orphans*.

9.2.4 Controlling the standard Linebreak Algorithm

PDFlib implements a sophisticated line-breaking algorithm. Table 9.2 lists Textflow options which control the line-breaking algorithm.

Line-breaking rules. When a word or other sequence of text surrounded by space characters doesn't fully fit into a line, it must be moved to the next line. In this situation the line-breaking algorithm decides after which characters a line break is possible.

For example, a formula such as $-12+235/8*45$ will never be broken, while the string *PDF-345+LIBRARY* may be broken to the next line at the minus character. If the text contains soft hyphen characters it can also be broken after such a character.

For parentheses and quotation marks it depends on whether we have an opening or closing character: opening parentheses and quotations marks do not offer any break

Fig. 9.21
Justified text with soft hyphen characters, using default settings and a wide fitbox

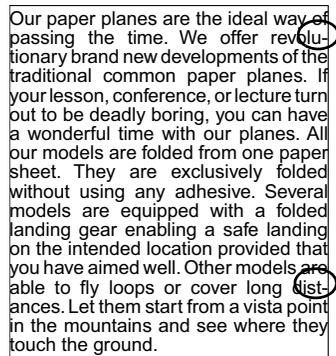
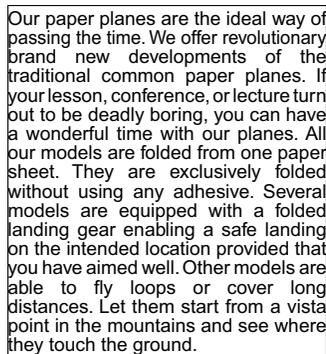


Fig. 9.22
Justified text without soft hyphens, using default settings and a wide fitbox.



opportunity. In order to find out whether a quotation mark starts or ends a sequence, pairs of quotation marks are examined.

Table 9.2 Options for controlling the line-breaking algorithm

option	explanation
adjust-method	<p>(Keyword) The method used to adjust a line when a text portion doesn't fit into a line after compressing or expanding the distance between words subject to the limits specified by the <code>minspacing</code> and <code>maxspacing</code> options. Default: <code>auto</code></p> <p>auto The following methods are applied in order: <code>shrink</code>, <code>spread</code>, <code>nofit</code>, <code>split</code>.</p> <p>clip Same as <code>nofit</code> (see below), except that the long part at the right edge of the fitbox (taking into account the <code>rightindent</code> option) will be clipped.</p> <p>nofit The last word will be moved to the next line provided the remaining (short) line will not be shorter than the percentage specified in the <code>nofitlimit</code> option. Even justified paragraphs will look slightly ragged in this case.</p> <p>shrink If a word doesn't fit in the line the text will be compressed subject to the <code>shrinklimit</code> option until the word fits. If it still doesn't fit the <code>nofit</code> method will be applied.</p> <p>split The last word will not be moved to the next line, but will forcefully be hyphenated. For text fonts a hyphen character will be inserted, but not for symbol fonts.</p> <p>spread The last word will be moved to the next line and the remaining (short) line will be justified by increasing the distance between characters in a word, subject to the <code>spreadlimit</code> option. If justification still cannot be achieved the <code>nofit</code> method will be applied.</p>
advanced-linebreak	<p>(Boolean) Enable the advanced line breaking algorithm which is required for complex scripts. This is required for linebreaking in scripts which do not use space characters for designating word boundaries, e.g. Thai. The options <code>locale</code> and <code>script</code> will be honored. Default: <code>false</code></p>
avoidbreak	<p>(Boolean) If <code>true</code>, avoid any line breaks until <code>avoidbreak</code> is reset to <code>false</code>. Default: <code>false</code></p>
charclass	<p>(List of pairs, where the first element in each pair is a keyword, and the second element is either a unichar or a list of unichars) The specified unichars will be classified by the specified keyword to determine the line breaking behavior of those character(s):</p> <p>letter behave like a letter (e.g. <code>a B</code>)</p> <p>punct behave like a punctuation character (e.g. <code>+ / ; :</code>)</p> <p>open behave like an open parenthesis (e.g. <code>[</code>)</p> <p>close behave like a close parenthesis (e.g. <code>]</code>)</p> <p>default reset all character classes to PDFlib's builtin defaults</p> <p>Example: <code>charclass={ close » open « letter {/ : =} punct & }</code></p>
hyphenchar	<p>(Unichar or keyword) Unicode value of the character which replaces a soft hyphen at line breaks. The value <code>o</code> and the keyword <code>none</code> completely suppress hyphens. Default: <code>U+00AD (SOFT HYPHEN)</code> if available in the font, <code>U+002D (HYPHEN-MINUS)</code> otherwise</p>
locale	<p>(Keyword) The locale which will be used for localized linebreaking methods if <code>advancedlinebreak= true</code>. The keywords consists of one or more components, where the optional components are separated by an underscore character <code>'_'</code> (the syntax slightly differs from NLS/POSIX locale IDs):</p> <ul style="list-style-type: none"> ▶ A required two- or three-letter lowercase language code according to ISO 639-2, e.g. <code>en</code> (English), <code>de</code> (German), <code>ja</code> (Japanese). This differs from the <code>language</code> option. ▶ An optional four-letter script code according to ISO 15924, e.g. <code>Hira</code> (Hiragana), <code>Hebr</code> (Hebrew), <code>Arab</code> (Arabic), <code>Thai</code> (Thai). ▶ An optional two-letter uppercase country code according to ISO 3166, e.g. <code>DE</code> (Germany), <code>CH</code> (Switzerland), <code>GB</code> (United Kingdom) <p>The keyword <code>_none</code> specifies that no locale-specific processing will be done.</p> <p>Specifying a locale is required for advanced line breaking for some scripts, e.g. Thai. Default: <code>_none</code></p> <p>Examples: <code>Thai</code>, <code>de_DE</code>, <code>en_US</code>, <code>en_GB</code></p>

Table 9.2 Options for controlling the line-breaking algorithm

option	explanation
maxspacing minspacing	(Float or percentage) Specifies the maximum or minimum distance between words (in user coordinates, or as a percentage of the width of the space character). The calculated word spacing is limited by the provided values (but the wordspacing option will still be added). Defaults: minspacing=50%, maxspacing=500%
nofitlimit	(Float or percentage) Lower limit for the length of a line with the nofit method (in user coordinates or as a percentage of the width of the fitbox). Default: 75%.
shrinklimit	(Percentage) Lower limit for compressing text with the shrink method; the calculated shrinking factor is limited by the provided value, but will be multiplied with the value of the horzscaling option. Default: 85%
spreadlimit	(Float or percentage) Upper limit for the distance between two characters for the spread method (in user coordinates or as a percentage of the font size); the calculated character distance will be added to the value of the charspacing option. Default: 0

An inline option list generally does not create a line break opportunity in order to allow option changes within words. However, when an option list is surrounded by space characters there is a line break opportunity at the beginning of the option list. If a line break occurs at the option list and *alignment=justify*, the spaces preceding the option list will be discarded. The spaces after the option list will be retained, and will appear at the beginning of the next line.

Preventing linebreaks. You can use the *charclass* option to prevent Textflow from breaking a line after specific characters. The following option prevents line breaks immediately after the / character:

```
charclass={letter /}
```

The following option prevents line breaks between the degree character »°« and the »C« in text like »20 °C«; the degree character is treated like an opening parenthesis which doesn't stand alone at the end of a line:

```
charclass={open °}
```

In order to prevent a sequence of text from being broken across lines you can bracket it with *avoidbreak...noavoidbreak*.

Cookbook A full code sample can be found in the *Cookbook* topic [textflow/avoid_linebreaking](#).

Formatting CJK text. The textflow engine is prepared to deal with CJK text, and properly treats CJK characters as ideographic glyphs as per the Unicode standard. As a result, CJK text will never be hyphenated. For improved formatting the following options are recommended when using Textflow with CJK text; they will disable hyphenation for inserted Latin text and create evenly spaced text output:

```
hyphenchar=none
alignment=justify
shrinklimit=100%
spreadlimit=100%
```

Vertical writing mode is not supported in Textflow.

Our paper planes are the ideal way of passing the time. We offer revolutionary brand new developments of the traditional common paper planes. If your lesson, conference, or lecture turn out to be deadly boring, you can have a wonderful time with our planes. All

Fig. 9.23 Justified text in a narrow fitbox with default settings

- decrease the distance between words (minspacing option)
- compress the line (shrink method, shrinklimit option)
- force hyphenation (split method)
- increase the distance between words (spread method, maxspacing option)

Justified text in a narrow fitbox. The narrower the fitbox, the more important are the options for controlling justified text. figure 9.23 demonstrates the results of the various methods for justifying text in a narrow fitbox. The option settings in figure 9.23 are basically okay, with the exception of *maxspacing* which provides a rather large distance between words. However, it is recommended to keep this for narrow fitboxes since otherwise the ugly forced hyphenation caused by the *split* method will occur more often.

If the fitbox is so narrow that occasionally forced hyphenations occur, you should consider inserting soft hyphens, or modify the options which control justified text.

Option shrinklimit for justified text. The most visually pleasing solution is to reduce the *shrinklimit* option which specifies a lower limit for the shrinking factor applied by the *shrink* method. figure 9.24a shows how to avoid forced hyphenation by compressing text down to *shrinklimit=50%*.

Fig. 9.24 Options for justified text in a narrow fitbox

	Generated output	Option list for <code>PDF_fit_textflow()</code>
a)	<p>passing the time. We offer revolutionary brand new developments of the traditional common paper planes. If your lesson, conference, or lecture turn out to</p>	<code>alignment=justify shrinklimit=50%</code>
b)	<p>Our paper planes are the ideal way of passing the time. We offer revolutionary brand new developments of the</p>	<code>alignment=justify spreadlimit=5</code>
c)	<p>ments of the traditional common paper planes. If your lesson, conference, or lecture turn out to be deadly boring, you can have</p>	<code>alignment=justify nofitlimit=50</code>

Option spreadlimit for justified text. Expanding text, which is achieved by the *spread* method and controlled by the *spreadlimit* option, is another method for controlling line

breaks. This unpleasing method should be rarely used, however. figure 9.24b demonstrates a very large maximum character distance of 5 units using `spreadlimit=5`.

Option `nofitlimit` for justified text. The `nofitlimit` option controls how small a line can get when the `nofit` method is applied. Reducing the default value of 75% is preferable to forced hyphenation when the fitbox is very narrow. figure 9.24c shows the generated text output with a minimum text width of 50%.

9.2.5 Advanced script-specific Line Breaking

PDFlib implements an additional line breaking algorithm on top of the standard line breaking algorithm. This advanced line breaking algorithm is required for some scripts, and improves line breaking behavior for some other script/locale combinations even if it is not required. It can be enabled with the `advancedlinebreak` option. Since line breaking depends on the language of the text, the advanced line breaking algorithm honors the `script` option (see Table 7.3) and the `locale` option (see PDFlib API Reference). Advanced line breaking determines proper line break opportunities in the following situations:

- ▶ For scripts in which line breaking does not rely on the presence of space characters in the text, e.g. Thai. The following Textflow option list enables advanced line breaking for Thai:

```
<advancedlinebreak script=thai locale=tha>
```

- ▶ In script/locale combinations which require specific treatment of certain punctuation characters, e.g. the « and » guillemet characters used as quotation marks in French text. The following Textflow option list enables advanced line breaking for French text. As a result, the guillemet characters surrounding a word will not be split apart from the word at the end of a line:

```
<advancedlinebreak script=latn locale=fr>
```

Note that the `locale` Textflow option is different from the `language` text option (see Table 7.4): although the `locale` option can start with the same three-letter language identifier, it can optionally contain one or two additional parts. However, these will rarely be required for PDFlib.

9.2.6 Wrapping Text around Paths and Images

The wrapping feature can be used to fill arbitrary shapes with text or wrap text around a path. By means of matchboxes, explicit rectangles/polygons/circles/curves or path objects you can specify wrapping areas for the Textflow. If an image contains an integrated clipping path you can wrap text around the image clipping path automatically.

Wrapping text around an image with matchbox. In the first example we will place an image within the Textflow and run the text around the whole image. First the image is loaded and placed in the box at the desired position. To refer to the image by name later, define a matchbox called `img` when fitting the image, and specify a margin of 5 units with the option list `matchbox={name=img margin=-5}` as follows:

```
result = p.fit_image(image, 50, 35,  
    "boxsize={80 46} fitmethod=meet position=center matchbox={name=img margin=-5}");
```

The Textflow is added. Then we place it using the *wrap* option with the image's matchbox *img* as the area to run around as follows (see figure 9.25):

```
result = p.fit_textflow(textflow, left_x, left_y, right_x, right_y,
    "wrap={usematchboxes={{img}}}");
```

Before placing the text you can fit more images using the same matchbox name. In this case the text will run around all images.

Cookbook A full code sample can be found in the Cookbook topic `textflow/wrap_text_around_images`.

Wrapping text around arbitrary paths. You can create a path object (see Section 3.2.3, »Direct Paths«, page 62) and use it as a wrap shape. The following fragment constructs a path object with a simple shape (a circle) and supplies it to the *wrap* option of `PDF_fit_textflow()`. The reference point for placing the path is expressed as percentage of the fitbox's width and height:

```
path = p.add_path_point( -1,  0, 100, "move", "");
path = p.add_path_point(path, 200, 100, "control", "");
path = p.add_path_point(path,  0, 100, "circular", "");

result = p.fit_textflow(tf, llx1, lly1, urx1, ury1,
    "wrap={paths={" +
        "{path=" + path + " refpoint={100% 50%} }" +
        "}}");

p.delete_path(path);
```

Use the *inversefill* option to wrap text inside the path instead of wrapping the text around the path (i.e. the path serves as text container instead of creating a hole in the Textflow):

```
result = p.fit_textflow(tf, llx1, lly1, urx1, ury1,
    "wrap={inversefill paths={" +
        "{path=" + path + " refpoint={100% 50%} }" +
        "}}");
```

Wrapping text around an image clipping path. TIFF and JPEG images can contain an integrated clipping path. The path must have been created in an image processing application and will be evaluated by PDFlib. If a default clipping path is found in the image it will be used, but you can specify any other clipping path in the image with the *clippingpathname* option of `PDF_load_image()`. If the image has been loaded with a clip-

Fig. 9.25
Wrapping text around an image with matchbox

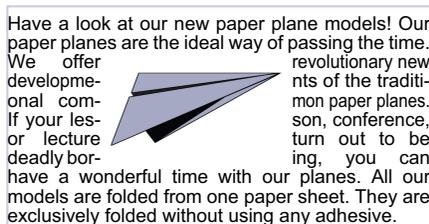
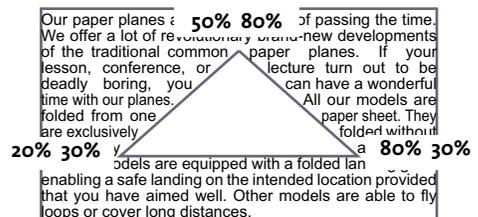


Fig. 9.26
Wrapping text around a triangular shape



ping path you can extract the path and supply it to the *wrap* option *PDF_fit_textflow()* as above. We also supply the *scale* option to enlarge the imported image clipping path:

```
image = p.load_image("auto", "image.tif", "clippingpathname={path 1}");

/* Create a path object from the image's clipping path */
path = (int) p.info_image(image, "clippingpath", "");
if (path == -1)
    throw new Exception("Error: clipping path not found!");

result = p.fit_textflow(tf, llx1, lly1, urx1, ury1,
    "wrap={paths={{path=" + path + " refpoint={50% 50%} scale=2}}});

p.delete_path(path);
```

Placing an image and wrapping text around it. While the previous section used only the clipping path of an image (but not the image itself), let's now place the image inside the fitbox of the Textflow and wrap the text around it. In order to achieve this we must again load the image with the *clippingpathname* option and place it on the page with *PDF_fit_image()*. In order to create the proper path object for wrapping the Textflow we call *PDF_info_image()* with the same option list as *PDF_fit_image()*. This ensures that the same transformations (e.g. scaling) are applied when the clipping path is retrieved and when the image is actually placed. Finally, the reference point (the *x/y* parameters of *PDF_fit_image()*) must be supplied to the *refpoint* suboption of the *paths* suboption of the *wrap* option:

```
image = p.load_image("auto", "image.tif", "clippingpathname={path 1}");

/* Place the image on the page with some fitting options */
String imageoptlist = "scale=2";
p.fit_image(image, x, y, imageoptlist);

/* Create a path object from the image's clipping path, using the same option list */
path = (int) p.info_image(image, "clippingpath", imageoptlist);
if (path == -1)
    throw new Exception("Error: clipping path not found!");

result = p.fit_textflow(tf, llx1, lly1, urx1, ury1,
    "wrap={paths={{path=" + path + " refpoint={ " + x + " " + y + " } }}}});

p.delete_path(path);
```

You can supply the same *wrap* option in multiple calls to *PDF_fit_textflow()*. This is useful if the placed image overlaps multiple Textflow fitboxes, e.g. for multi-column text.

Wrapping text around non-rectangular shapes. As an alternative to creating a path object as wrap shape you can specify path elements directly in Textflow options.

In addition to wrapping text around a rectangle specified by a matchbox you can define arbitrary graphical elements as wrapping shapes. For example, the following option list will wrap the text around a triangular shape (see figure 9.26):

```
wrap={ polygons={ {50% 80% 20% 30% 80% 30% 50% 80%} } }
```

Fig. 9.27
Filling a rhombus shape with text

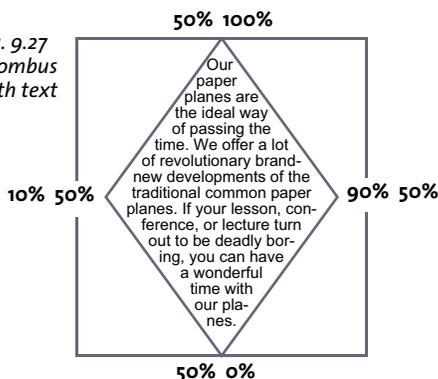
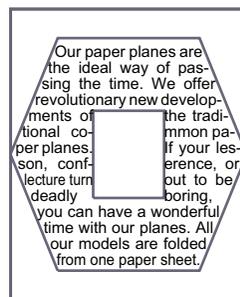


Fig. 9.28
Filling overlapping shapes



Note that the `showborder=true` option has been used to illustrate the margins of the shapes. The `wrap` option can contain multiple shapes. The following option list will wrap the text around two triangle shapes:

```
wrap={ polygons={ {50% 80% 20% 30% 80% 30% 50% 80%}
                  {20% 90% 10% 70% 30% 70% 20% 90%} } }
```

Instead of percentages (relative coordinates within the fitbox) absolute coordinates on the page can be used.

Note It is recommended to set `fixedleading=true` when using shapes with segments which are neither horizontally nor vertically oriented.

Cookbook A full code sample can be found in the Cookbook topic `textflow/wrap_text_around_polygons`.

Filling arbitrary shapes. The wrap feature can also be used to place the contents of a Textflow in arbitrarily shaped areas. This is achieved with the `addfitbox` or `inversefill` sub-options of the `wrap` option. Instead of wrapping the text around the specified shapes the text will be placed within one or more shapes. The following option list can be used to flow text into a rhombus shape, where the coordinates are provided as percentages of the fitbox rectangle (see figure 9.27):

```
wrap={ addfitbox polygons={ {50% 100% 10% 50% 50% 0% 90% 50% 50% 100%} } }
```

The option `showborder=true` has been used again to illustrate the margins of the shape. Without the `addfitbox` option the rhombus shape will remain empty and the text will be wrapped around it.

Unless the top of the wrap shape coincides with the top of the fitbox, the formatting process doesn't necessarily start at the top of the wrap box. Instead, the formatter attempts to place the text vertically at a line which is determined from the start of the fitbox (not at top of the wrap shape), but such that the text is contained inside the wrap shape. This may require reduced font size according to the selected `fitmethod`. In a sense the lines within the fitbox form a grid and the wrap shape is used as window into this grid. If the top of the wrap shape doesn't coincide with the horizontal position of a line, empty space may occur at the top of the wrap shape above the first text output. This is a result of the grid-based formatting described above.

Filling overlapping shapes. In the next example we will fill a shape comprising two overlapping polygons, namely a hexagon with a rectangle inside. Using the *addfitbox* option the fitbox itself will be excluded from being filled, and the polygons in the subsequent list will be filled except in the overlapping area (see figure 9.28):

```
wrap={ addfitbox polygons=
  { {20% 10% 80% 10% 100% 50% 80% 90% 20% 90% 0% 50% 20% 10%}
    {35% 35% 65% 35% 65% 65% 35% 65% 35% 35% } } }
```

Without the *addfitbox* option you will get the opposite effect: the previously filled area will remain empty, and the previously empty areas will be filled with text.

Cookbook A full code sample can be found in the *Cookbook* topic `textflow/fill_polygons_with_text`.

9.3 Table Formatting

The table formatting feature can be used to automatically format complex tables. Table cells may contain single- or multi-line text, images, SVG graphics or PDF pages. Tables are not restricted to a single fitbox, but can span multiple pages.

Cookbook Code samples regarding table issues can be found in the *table* category of the *PDFlib Cookbook*.

General aspects of a table. The description of the table formatter is based on the following concepts and terms (see figure 9.29):

- ▶ A *table* is a virtual object with a rectangular outline. It is comprised of horizontal *rows* and vertical *columns*.
- ▶ A *simple cell* is a rectangular area within a table, defined as the intersection of a row and a column. A *spanning cell* spans more than one column, more than one row, or both. The term *cell* will be used to designate both simple and spanning cells.
- ▶ The complete table may fit into one fitbox, or several fitboxes may be required. The rows of the table which are placed in one fitbox constitute a *table instance*. Each call to *PDF_fit_table()* will place one *table instance* in one fitbox (see Section 9.3.5, »Table Instances«, page 235).
- ▶ The *header* or *footer* is a group of one or more rows at the beginning or end of the table which are repeated at the top or bottom of each table instance. Rows which are neither part of the header nor footer are called *body rows*.

- ▶ An optional caption (not shown in figure 9.29) is an auxiliary element which can be used to place a description of the table. It can be placed on any side of the table.

Our Paper Plane Models		
1 Giant Wing		
Material	Offset print paper 220g/sqm	
Benefit	It is amazingly robust and can even do aerobatics. But it is best suited to gliding.	
2 Long Distance Glider		
Material	Drawing paper 180g/sqm	
Benefit	With this paper rocket you can send all your messages even when sitting in the cinema pretty near the back.	
3 Cone Head Rocket		
Material	Kent paper 200g/sqm	
Benefit	This paper arrow can be thrown with big swing. It stays in the air a long time.	

Annotations in the diagram:

- Header: Points to the top row of the table.
- Simple cell: Points to the 'Material' cell in the first model's row.
- Cell containing image and text line: Points to the 'Benefit' cell in the first model's row, which contains an image of a paper plane and the text 'Amazingly robust!'.
- Cell spanning three columns: Points to the header row.
- Cell spanning three rows: Points to the 'Benefit' cell in the second model's row, which contains an image of a paper plane and the text 'With big swing!'.
- Cell containing Textflow: Points to the 'Benefit' cell in the second model's row.
- Footer: Points to the bottom row of the table.

Fig. 9.29 Sample table

As an example, all aspects of creating the table in figure 9.29 will be explained. A complete description of the table formatting options can be found in the *PDFlib API Reference*. Creating a table starts by defining the contents and visual properties of each table cell with `PDF_add_table_cell()`. Then you place the table using one or more calls to `PDF_fit_table()`.

When placing the table the size of its fitbox and the ruling and shading of table rows or columns can be specified. Use the Matchbox feature for details such as cell-specific shading (see Section 9.4, »Matchboxes«, page 241, for more information).

In this section the most important options for defining the table cells and fitting the table will be discussed. All examples demonstrate the relevant calls of `PDF_add_table_cell()` and `PDF_fit_table()` only, assuming that the required font has already been loaded.

Note Table processing is independent from the current graphics state. Table cells can be defined in document scope while the actual table placement must be done in page or pattern/template/glyph scope.

Cookbook A full code sample can be found in the Cookbook topic table/starter_table.

Analyzing table-related formatting problems. Depending on the number of cells and the table formatting options the results of PDFlib's table formatter may sometimes not match your expectations. In almost all cases this can be rectified with suitable options. However, it may be difficult to identify a problematic cells or group of cells with the wrong options. In order to facilitate debugging of table-related formatting problems PDFlib offers the following options in `PDF_fit_table()`:

- ▶ The option *showcells* visualizes the border of each inner cell box. If the method is called in page scope and PDF/A mode is not active, an annotation with details about the cell contents is placed in the center of each table cell.
- ▶ If the option *debugshow* is *true*, all errors for tables which are too high, too wide, or where the cells get too small are suppressed and are logged instead. The resulting table instance is created as a debugging aid although the table is damaged.
- ▶ If the option *showgrid* is *true*, the vertical and horizontal boundary of all columns and rows are stroked, i.e. the underlying tabular grid is visualized.

9.3.1 Placing a Simple Table

Before we describe the table concepts in more detail, we will demonstrate a simple example for creating a table. The table contains six cells which are arranged in three rows and two columns. Four cells contain text lines, and one cell contains a multi-line Textflow. All cell contents are horizontally aligned to the left, and vertically aligned to the center with a margin of 4 points.

To create the table we first prepare the option list for the text line cells by defining the required options *font* and *fontsize* and a position of *{left center}* in the *fittextline* sub-option list. In addition, we define cell margins of 4 points. Then we add the text line cells one after the other in their respective column and row, with the actual text supplied directly in the call to *PDF_add_table_cell()*.

In the next step we create a Textflow, use the Textflow handle to assemble the option list for the Textflow table cell, and add that cell to the table.

Finally we place the table with *PDF_fit_table()* while visualizing the table frame and cell borders with black lines. Since we didn't supply any column widths, they will be calculated automatically from the supplied text lines plus the margins.

Cookbook A full code sample can be found in the *Cookbook* topic [table/vertical_text_alignment](#).

The following code fragment shows how to create the simple table. The result is shown in figure 9.30a.

```
/* Text for filling a table cell with multi-line Textflow */
String tf_text = "It is amazingly robust and can even do aerobatics. " +
    "But it is best suited to gliding.";

/* Define the column widths of the first and the second column */
int c1 = 80, c2 = 120;

/* Define the lower left and upper right corners of the table instance (fitbox) */
double llx=100, lly=500, urx=300, ury=600;

/* bail out on errors */
p.set_option("errorpolicy=exception");

/* Load the font */
font = p.load_font("Helvetica", "unicode", "");

/* Define the option list for the text line cells placed in the first column */
optlist = "fittextline={position={left center} font=" + font + " fontsize=8} margin=4" +
    colwidth=" + c1;

/* Add a text line cell in column 1 row 1 */
tbl = p.add_table_cell(tbl, 1, 1, "Our Paper Planes", optlist);
```

```

/* Add a text line cell in column 1 row 2 */
tbl = p.add_table_cell(tbl, 1, 2, "Material", optlist);

/* Add a text line cell in column 1 row 3 */
tbl = p.add_table_cell(tbl, 1, 3, "Benefit", optlist);

/* Define the option list for a text line placed in the second column */
optlist = "fittextline={position={left center} font=" + font + " fontsize=8} " +
"colwidth=" + c2 + " margin=4";

/* Add a text line cell in column 2 row 2 */
tbl = p.add_table_cell(tbl, 2, 2, "Offset print paper 220g/sqm", optlist);

/* Add a Textflow */
optlist = "font=" + font + " fontsize=8 leading=110%";
tf = p.add_textflow(-1, tf_text, optlist);

/* Define the option list for the Textflow cell using the handle retrieved above */
optlist = "textflow=" + tf + " margin=4 colwidth=" + c2;

/* Add the Textflow table cell in column 2 row 3 */
tbl = p.add_table_cell(tbl, 2, 3, "", optlist);

p.begin_page_ext(0, 0, "width=200 height=100");

/* Define the option list for fitting the table with table frame and cell ruling */
optlist = "stroke={{line=frame linewidth=0.8} {line=other linewidth=0.3}}";

/* Place the table instance */
result = p.fit_table(tbl, llx, lly, urx, ury, optlist);

/* Check the result; "_stop" means all is ok. */
if (!result.equals("_stop")) {
    if (result.equals("_error"))
        throw new Exception("Error: " + p.get_errmsg());
    else {
        /* Any other return value requires dedicated code to deal with */
    }
}
p.end_page_ext("");

/* This also deletes Textflow handles used in the table */
p.delete_table(tbl, "");

```

Fine-tuning the vertical alignment of cell contents. When we vertically center contents of various types in the table cells, they will be positioned with varying distance from the borders. In figure 9.30a, the four text line cells have been placed with the following option list:

```

optlist = "fittextline={position={left center} font=" + font +
" fontsize=8} colwidth=80 margin=4";

```

The Textflow cell is added without any special options. Since we vertically centered the text lines, the *Benefit* line will move down with the height of the Textflow.

Fig. 9.30 Aligning text lines and Textflow in table cells

Generated output		
a)	Our Paper Planes	
	Material	Offset print paper 220g/sqm
	Benefit	It is amazingly robust and can even do aerobatics. But it is best suited to gliding.
b)	Our Paper Planes	
	Material	Offset print paper 220g/sqm
	Benefit	It is amazingly robust and can even do aerobatics. But it is best suited to gliding.

As shown in figure 9.30b, we want all cell contents to have the same *vertical* distance from the cell borders regardless of whether they are Textflows or text lines. To accomplish this we first prepare the option list for the text lines. We define a fixed row height of 14 points, and the position of the text line to be on the top left with a margin of 4 points.

The `fontsize=8` option which we supplied before doesn't exactly represent the letter height but adds some space below and above. However, the height of an uppercase letter is exactly represented by the `capheight` value of the font. For this reason we use `fontsize={capheight=6}` which will approximately result in a font size of 8 points and (along with `margin=4`), will sum up to an overall height of 14 points corresponding to the `rowheight` option. The complete option list of `PDF_add_table_cell()` for our text line cells looks as follows:

```
/* option list for the text line cells */
optlist = "fittextline={position={left top} font=" + font +
         " " fontsize={capheight=6} rowheight=14 colwidth=80 margin=4";
```

To add the Textflow we use `fontsize={capheight=6}` which will approximately result in a font size of 8 points and (along with `margin=4`), will sum up to an overall height of 14 points as for the text lines above.

```
/* option list for adding the Textflow */
optlist = "font=" + font + " " fontsize={capheight=6} leading=110%";
```

In addition, we want the baseline of the *Benefit* text aligned with the first line of the Textflow. At the same time, the *Benefit* text should have the same distance from the top cell border as the *Material* text. To avoid any space from the top we add the Textflow cell using `fittextflow={firstlinedist=capheight}`. Then we add a margin of 4 points, the same as for the text lines:

```
/* option list for adding the Textflow cell */
optlist = "textflow=" + tf + " " fittextflow={firstlinedist=capheight} "
         " colwidth=120 margin=4";
```

9.3.2 Contents of a Table Cell

When adding cells to a table with `PDF_add_table_cell()`, you can specify various kinds of cell contents. Table cells can contain one or more content types at the same time. Additional ruling and shading is available, as well as matchboxes which can be used for placing additional content in a table cell.

For example, the cells of the paper plane table contain the elements illustrated in figure 9.31.

Text line	
Text line	
Text line	Text line
Text line	Textflow

Fig. 9.31
Contents of the
table cells

Single-line text with Textlines. The text is supplied in the `text` parameter of `PDF_add_table_cell()`. In the `fittextline` option list all formatting options of `PDF_fit_textline()` can be specified. The default fit method is `fitmethod=nofit`. The cell is enlarged if the text doesn't completely fit into the cell. To avoid this use `fitmethod=auto` to shrink the text, subject to the `shrinklimit` option. If no row height is specified, the formatter assumes twice the text size as height of the table cell (more precisely: twice the `boxheight`, which has the default value `{capheight none}` unless specified otherwise). The same applies to the row width for rotated text.

Use the `fillcolor` option in the `fittextline` option list to apply color to the text.

Multi-line text with Textflow. The Textflow must have been prepared outside the table functions and created with `PDF_create_textflow()` or `PDF_add_textflow()` before calling `PDF_add_table_cell()`. The Textflow handle is supplied in the `textflow` option. In the `fittextflow` option all formatting options of `PDF_fit_textflow()` can be specified.

The default fit method is `fitmethod=clip`: first it is attempted to completely fit the text into the cell. If the cell is not large enough its height is increased. If the text does not fit anyway it is clipped at the bottom. To avoid this use `fitmethod=auto` to shrink the text subject to the `minfontsize` option.

If the cell is too narrow the Textflow could be forced to split single words at undesired positions. If the `avoidwordsplitting` option is `true` the cell width will be enlarged until no more word splitting occurs.

Images and templates. Images must be loaded with `PDF_load_image()` before calling `PDF_add_table_cell()`. Templates must be created with `PDF_begin_template_ext()`. The image or template handle is supplied in the `image` option. In the `fitimage` option all formatting options of `PDF_fit_image()` can be specified. The default fit method is `fitmethod=meet`: the image/template is placed completely inside the cell without distorting its aspect ratio. The cell size is not changed due to the size of the image/template.

Vector graphics. Graphics must be loaded with `PDF_load_graphics()` before calling `PDF_add_table_cell()`. The graphics handle is supplied in the `graphics` option. In the `fitgraphics` option all formatting options of `PDF_fit_graphics()` can be specified. The default fit method is `fitmethod=meet`: the graphics is placed completely inside the cell without distorting its aspect ratio. The cell size is not changed due to the size of the graphics.

Pages from an imported PDF document. The PDI page must have been opened with `PDF_open_pdi_page()` before calling `PDF_add_table_cell()`. The PDI page handle is supplied in the `pdipage` option. In the `fitpdipage` option all formatting options of `PDF_fit_pdi_page()` can be specified. The default fit method is `fitmethod=meet`: the PDI page is placed completely inside the cell without distorting its aspect ratio. The cell size is not changed due to the size of the PDI page.

Path objects. Path objects must have been created with `PDF_add_path_point()` before calling `PDF_add_table_cell()`. The path handle is supplied in the `path` option. In the `fitpath` option all formatting options of `PDF_draw_path()` can be specified. The bounding box of the path is placed in the table cell. The lower left corner of the inner cell box is used as reference point for placing the path.

Annotations. Annotations in table cells can be created with the `annotationtype` option of `PDF_add_table_cell()` which corresponds to the `type` parameter of `PDF_create_annotation()` (but this method does not have to be called). In the `fitannotation` option all options of `PDF_create_annotation()` can be specified. The cell box is used as annotation rectangle.

Form fields. Form fields in table cells can be created with the `fieldname` and `fieldtype` options of `PDF_add_table_cell()` which correspond to the `name` and `type` parameters of `PDF_create_field()` (but this method does not have to be called). In the `fitfield` option all options of `PDF_create_field()` can be specified. The cell box is used as field rectangle.

Positioning cell contents in the inner cell box. By default, cell contents are positioned with respect to the cell box. The `margin` options of `PDF_add_table_cell()` can be used to specify some distance from the cell borders. The resulting rectangle is called the *inner cell box*. If any margin is defined, the cell contents are placed with respect to the inner cell box (see figure 9.32). If no margins are defined, the inner cell box is identical to the cell box.

In addition, cell contents may be subject to further options supplied in the content-specific fit options as described in section Section 9.3.4, «Mixed Table Contents», page 233.

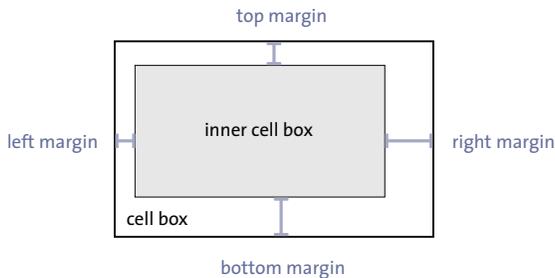


Fig. 9.32
Fitting contents in
the inner cell box

9.3.3 Table and Column Widths

When adding a cell to the table, you define the number of columns and/or rows spanned by the cell with the *colspan* and *rowspan* options. By default, a cell spans one column and one row. The total number of columns and rows in the table is implicitly increased by the respective values when adding a cell. figure 9.33 shows an example of a table containing three columns and four rows.

row 1	cell spanning three columns		
row 2	cell spanning two columns		cell spanning three rows
row 3	simple cell	simple cell	
row 4	simple cell	simple cell	
	column 1	column 2	column 3

Fig. 9.33
Simple cells and cells spanning several rows or columns

Furthermore you can explicitly supply the width of the first column spanned by the cell with the *colwidth* option. Keep in mind that the supplied *colwidth* affects only the first column, not all columns in the *colspan* group. By supplying each cell with a defined first column width all width values implicitly add up to the total table width. figure 9.34 shows an example.

colspan=3 colwidth=50			rowspan=3 colwidth=90
colspan=2 colwidth=50			
colspan=1 colwidth=50	colspan=1 colwidth=100		
colspan=1 colwidth=50	colspan=1 colwidth=100		
50 100 90			
total table width of 240			

Fig. 9.34
Column widths define the total table width.

Alternatively, you can specify the column widths as percentages if appropriate. In this case the percentages refer to the width of the table's fitbox. Either none or all column widths must be supplied as percentages.

If some columns are combined to a column scaling group with the *colscalegroup* option of *PDF_add_table_cell()*, their widths are adjusted to the widest column in the group (see figure 9.35),

	column scaling group			
	Max. Load	Range	Weight	Speed
Giant Wing	12g	18m	14g	8m/s
Long Distance Glider	5g	30m	11.2g	5m/s
Cone Head Rocket	7g	7m	12.4g	6m/s

Fig. 9.35
The last four cells in the first row are in the same column scaling group. They will have the same widths.

If absolute coordinates are used (as opposed to percentages) and there are cells left without any column width defined, the missing widths are calculated as follows: First, for each cell containing a text line the actual width is calculated based on the column width or the text width (or the text height in case of rotated text). Then, the remaining table width is evenly distributed among the column widths which are still missing.

9.3.4 Mixed Table Contents

In the following sections we will create the sample table containing various kinds of contents as shown in figure 9.36 step by step.

Cookbook A full code sample can be found in the *Cookbook topic* table/mixed_table_contents.

As a prerequisite we need to load two fonts. We define the dimensions of the table's fit-box in terms of the coordinates of the lower left and upper right corners and specify the widths of the three table columns. Then, we start a new page with A4 size:

```
double llx = 100, lly = 500, urx = 360, ury = 600; // coordinates of the table

int c1 = 50, c2 = 120, c3 = 90; // widths of the three table columns

boldfont = p.load_font("Helvetica-Bold", "unicode", "");
normalfont = p.load_font("Helvetica", "unicode", "");

p.begin_page_ext(0, 0, "width=a4.width height=a4.height");
```

Step 1: Adding the first cell. We start with the first cell of our table. The cell will be placed in the first column of the first row and will span three columns. The first column has a width of 50 points. The text line is centered vertically and horizontally, with a margin of 4 points from all borders. The following code fragment shows how to add the first cell:

```
optlist = "fittextline={font=" + boldfont + " fontsize=12 position=center} " +
          "fillcolor=red margin=4 colspan=3 colwidth=" + c1;

tbl = p.add_table_cell(tbl, 1, 1, "Our Paper Plane Models", optlist);
```

Step 2: Adding one cell spanning two columns. In the next step we add the cell containing the text line *Giant Wing*. It will be placed in the first column of the second row and spans two columns. The first column has a width of 50 points. The row height is 14 points. The text line is positioned on the top left, with a margin of 4 points from all borders. We use `fontsize={capheight=6}` to get a unique vertical text alignment as described in »Fine-tuning the vertical alignment of cell contents«, page 228.

Since the *Giant Wing* heading cell doesn't cover a complete row but only two of three columns it cannot be filled with color using one of the row-based shading options. We apply the Matchbox feature instead to fill the rectangle covered by the cell with a gray background color. The Matchbox feature is discussed in detail in Section 9.4, »Matchboxes«, page 241. The following code fragment demonstrates how to add the *Giant Wing* heading cell:

```
optlist = "fittextline={position={left top} font=" + boldfont +
          " fontsize={capheight=6} rowheight=14 colwidth=" + c1 +
          " margin=4 colspan=2 matchbox={fillcolor={gray .92}}";
```

```
tbl = p.add_table_cell(tbl, 1, 2, "1 Giant Wing", optlist);
```

Fig. 9.36 Adding table cells with various contents step by step

Generated table		Generation steps								
<table border="1"> <thead> <tr> <th colspan="2">Our Paper Plane Models</th> </tr> </thead> <tbody> <tr> <td>1 Giant Wing</td> <td rowspan="3">  </td> </tr> <tr> <td>Material</td> <td>Offset print paper 220g/sqm</td> </tr> <tr> <td>Benefit</td> <td>It is amazingly robust and can even do aerobatics. But it is best suited to gliding.</td> </tr> </tbody> </table>		Our Paper Plane Models		1 Giant Wing		Material	Offset print paper 220g/sqm	Benefit	It is amazingly robust and can even do aerobatics. But it is best suited to gliding.	<p>Step 1: Add a cell spanning 3 columns</p> <p>Step 2: Add a cell spanning 2 columns</p> <p>Step 3: Add 3 more text line cells</p> <p>Step 4: Add the Textflow cell</p> <p>Step 5: Add the image cell with a text line</p> <p>Step 6: Fitting the table</p>
Our Paper Plane Models										
1 Giant Wing										
Material		Offset print paper 220g/sqm								
Benefit		It is amazingly robust and can even do aerobatics. But it is best suited to gliding.								

Step 3: Add three more Textline cells. The following code fragment adds the *Material*, *Benefit* and *Offset print paper...* cells. The *Offset print paper...* cell will start in the second column defining a column width of 120 points. The cell contents is positioned on the top left, with a margin of 4 points from all borders.

```
optlist = "fittextline={position={left top} font=" + normalfont +
          " fontsize={capheight=6}} rowheight=14 colwidth=" + c1 + " margin=4";

tbl = p.add_table_cell(tbl, 1, 3, "Material", optlist);
tbl = p.add_table_cell(tbl, 1, 4, "Benefit", optlist);

optlist = "fittextline={position={left top} font=" + normalfont +
          " fontsize={capheight=6}} rowheight=14 colwidth=" + c2 + " margin=4";

tbl = p.add_table_cell(tbl, 2, 3, "Offset print paper 220g/sqm", optlist);
```

Step 4: Add the Textflow cell. The following code fragment adds the *It is amazingly...* Textflow cell. To add a table cell containing a Textflow we first create the Textflow. We use *fontsize={capheight=6}* which will approximately result in a font size of 8 points and (along with *margin=4*), will sum up to an overall height of 14 points as for the text lines above.

```
tftext = "It is amazingly robust and can even do aerobatics. " +
         "But it is best suited to gliding.";

optlist = "font=" + normalfont + " fontsize={capheight=6} leading=110%";

tf = p.add_textflow(-1, tftext, optlist);
```

The retrieved Textflow handle will be used when adding the table cell. The first line of the Textflow should be aligned with the baseline of the *Benefit* text line. At the same time, the *Benefit* text should have the same distance from the top cell border as the *Material* text. Add the Textflow cell using *fittextflow={firstlinedist=capheight}* to avoid any space from the top. Then add a margin of 4 points, the same as for the text lines:

```
optlist = "textflow=" + tf + " fittextflow={firstlinedist=capheight} " +
          " colwidth=" + c2 + " margin=4";

tbl = p.add_table_cell(tbl, 2, 4, "", optlist);
```

Step 5: Add the image cell with a text line. In the fifth step we add a cell containing an image of the Giant Wing paper plane as well as the *Amazingly robust!* text line. The cell will start in the third column of the second row and spans three rows. The column width is 90 points. The cell margins are set to 4 points. For a first variant we place a TIFF image in the cell:

```
image = p.load_image("auto", "kraxi_logo.tif", "");
optlist = "fittextline={font=" + boldfont + " fontsize=9} image=" + image +
          " colwidth=" + c3 + " rowspan=3 margin=4";
tbl = p.add_table_cell(tbl, 3, 2, "Amazingly robust!", optlist);
```

Alternatively, you could import the image as a PDF page. Make sure that the PDI page is closed only after the call to `PDF_fit_table()`:

```
int doc = p.open_pdi("kraxi_logo.pdf", "", 0);
page = p.open_pdi_page(doc, pageno, "");
optlist = "fittextline={font=" + boldfont + " fontsize=9} pdipage=" + page +
          " colwidth=" + c3 + " rowspan=3 margin=4";
tbl = p.add_table_cell(tbl, 3, 2, "Amazingly robust!", optlist);
```

Step 6: Fit the table. In the last step we place the table with `PDF_fit_table()`. Using `header=1` the table header will include the first line. The `fill` option and the suboptions `area=header` and `fillcolor={rgb 0.8 0.8 0.8}` specify the header row(s) to be filled with the supplied color. Using the `stroke` option and the suboptions `line=frame linewidth=0.8` we define a ruling of the table frame with a line width of 0.8. Using `line=other linewidth=0.3` a ruling of all cells is specified with a line width of 0.3:

```
optlist = "header=1 fill={{area=header fillcolor={rgb 0.8 0.8 0.8}}}" +
          "stroke={{line=frame linewidth=0.8} {line=other linewidth=0.3}}";
result = p.fit_table(tbl, llx, lly, urx, ury, optlist);
if (result.equals("_error"))
    throw new Exception("Error: " + p.get_errmsg());
p.end_page_ext("");
```

9.3.5 Table Instances

The rows of the table which are placed in one fitbox comprise a table instance. One or more table instances may be required to represent the full table. Each call to `PDF_fit_table()` will place one table instance in one fitbox. The fitboxes can be placed on the same page, e.g. with a multi-column layout, or on several pages.

The table in figure 9.37 is spread over three pages. Each table instance is placed in one fitbox on one page. For each call to `PDF_fit_table()` the first row is defined as header and the last row is defined as footer.

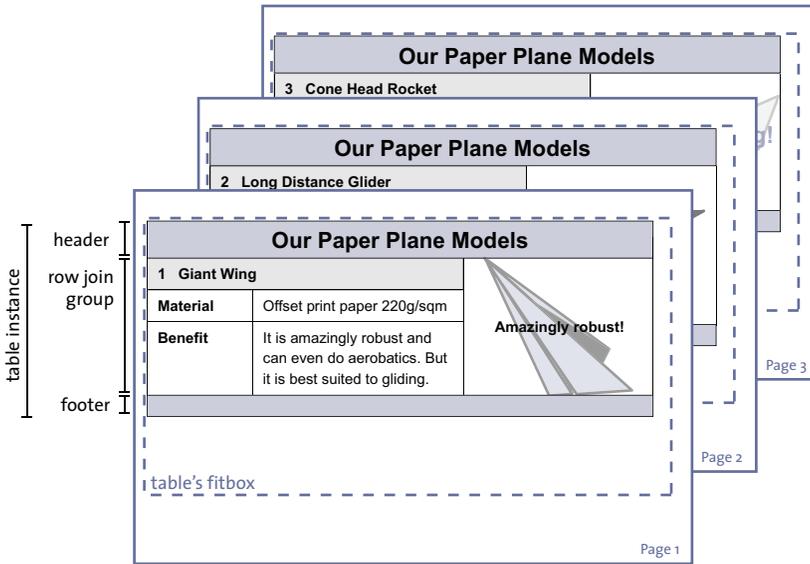


Fig. 9.37
Table broken into several table instances placed in one fitbox each.

The following code fragment shows the general loop for fitting table instances until the table has been placed completely. New pages are created as long as more table instances need to be placed.

```
do {
    /* Create a new page */
    p.begin_page_ext(0, 0, "width=a4.width height=a4.height");

    /* Use the first row as header and draw lines for all table cells */
    optlist = "header=1 stroke={{line=other}}";

    /* Place the table instance */
    result = p.fit_table(tbl, llx, lly, urx, ury, optlist);
    if (result.equals("_error"))
        throw new Exception("Error: " + p.get_errmsg());

    p.end_page_ext("");
} while (result.equals("_boxfull"));

/* Check the result; "_stop" means all is ok. */
if (!result.equals("_stop")) {
    if (result.equals("_error"))
        throw new Exception("Error: " + p.get_errmsg());
    else {
        /* Any other return value is a user exit caused by the "return" option;
        * this requires dedicated code to deal with. */
        throw new Exception ("User return found in Textflow");
    }
}
```

```
/* This will also delete Textflow handles used in the table */
p.delete_table(tbl, "");
```

Headers and footers. With the *header* and *footer* options of *PDF_fit_table()* you can define the number of initial or final table rows which will be placed at the top or bottom of a table instance. Using the *fill* option with *area=header* or *area=footer*, headers and footers can be individually filled with color. Header rows consist of the first *n* rows of the table definition and footer rows of the last *m* rows.

Headers and footers are specified per table instance in *PDF_fit_table()*. Consequently, they can differ among table instances: while some table instances include headers/footers, others can omit them, e.g. to specify a special row in the last table instance.

Joining rows. In order to ensure that a set of rows will be kept together in the same table instance, they can be assigned to the same row join group using the *rowjoingroup* option. The row join group contains multiple consecutive rows. The rows in the group are prevented from being separated into multiple table instances.

The rows of a cell spanning these rows don't constitute a join group automatically.

Splitting a cell. If the last rows spanned by a cell don't fit in the fitbox the cell is split. In case of an image, PDI page, SVG graphics or text line cell, the cell contents are repeated in the next table instance. In case of a Textflow cell the cell contents continues in the remaining rows of the cell.

figure 9.38 shows how the Textflow cell is split while the Textflow continues in the next row. In figure 9.39 an image cell is shown which is repeated in the first row of the next table instance.

table instance 1	1 Giant Wing		Our paper planes are the ideal way of passing the time. We offer revolutionary
	Material	Offset print paper 220g/sqm	
table instance 2	Benefit	It is amazingly robust and can even do aerobatics. But it is best suited to gliding.	new developments of the traditional common paper planes.

Fig. 9.38
Splitting a cell

Splitting a row. If the last body row doesn't completely fit into the table's fitbox it will usually not be split. This behavior is controlled by the *minrowheight* option of *PDF_fit_table()* with a default value of 100%. With this default setting the row is not split but completely placed in the next table instance.

If you decrease the *minrowheight* value the last body row is split with the given percentage of contents in the first instance and the remaining parts of the row are placed in the next instance.

figure 9.39 illustrates how the Textflow *It's amazingly robust...* is split and the Textflow is continued in the first body row of the next table instance. The image cell spanning several rows is split and the image is repeated. The *Benefit* cell is repeated as well.

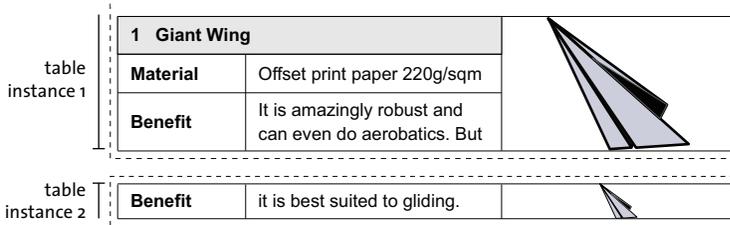


Fig. 9.39
Splitting a row

9.3.6 Table Formatting Algorithm

This section details the steps performed by the table formatter when placing a table. The description below applies to horizontal text. However, if you swap the terms »row height« and »column width« it also applies to vertical or rotated text.

In the first call to `PDF_fit_table()` the options `colwidth`, `rowheight`, `fittextline`, and `fittextflow` are examined for all cells, and the width and height of the full table is calculated based on column widths, row heights, Textline and Textflow contents and the width of the first fitbox. The height of the fitbox is assumed as infinite. The first table instance (i.e. the placement of the first part of the table in the first fitbox) is calculated according to the `fitmethod` option of `PDF_fit_table()`.

Calculate the height and width of table cells with Textlines. The table formatter initially determines the size of all table cells with Textlines which span table columns or rows without `colwidth` or `rowheight`. In order to achieve this, it calculates the width of the Textline and therefore the table cell according to the `fittextline` option. It assumes twice the text size as height of the table cell (more precisely: twice the box height, which has the default value `{capheight none}` unless specified otherwise). For vertical text the width of the widest character is used as cell width. For text oriented to west or east twice the text height is used as cell width.

The remaining width and height of the table cell is then distributed among all columns and rows spanned by the cell for which `colwidth` or `rowheight` hasn't been specified.

Calculate a tentative table size. In the next step the formatter calculates a tentative table width and height as the sum of all column widths and row heights, respectively. Column widths and row heights specified as percentages are converted to absolute values based on the width and height of the first fitbox. If there are still columns or rows without `colwidth` or `rowheight` the remaining space is evenly distributed until the tentative table size equals the first fitbox.

The `rowheightdefault` option can be used to completely fill the height of the fitbox (keywords `auto` and `distribute`) or to save space (keyword `minimum`). Explicitly specifying the height of a row with the `rowheight` option overrides `rowheightdefault`.

Enlarge cells which are too small. Now the formatter determines all inner cell boxes (see figure 9.32). If the combined margins are larger than the cell's width or height, the

cell box is suitably enlarged by evenly enlarging all columns and rows which belong to the cell.

Fit Textlines horizontally. The formatter attempts to increase the width of all cells with Textlines so that the Textline fits into the cell without reducing the font size. If this is not possible, the Textline is automatically placed with *fitmethod=auto*. This guarantees that the Textline does not extend beyond the inner cell box. You can prevent the cell width from being increased by setting *fitmethod=auto* in the *fittextline* option.

You can use the *colscalegroup* option to make sure that all columns which belong to the same column scaling group are scaled to equal widths, i.e. their widths are unified and adjusted to the widest column in the group (see figure 9.35).

Avoid forced hyphenation. If the calculated table width is smaller than the fitbox the formatter tries to increase the width of a Textflow cell so that the text fits without forced hyphenation. This can be avoided with the option *avoidwordsplitting=false*. The widths of such cells are increased until the table width equals the width of the fitbox.

You can query the difference between table width and fitbox width with the *horbox-gap* key of *PDF_info_table()*.

Fit text vertically. The formatter attempts to increase the height of all Textline and Textflow cells so that the Textline or Textflow fits into the inner cell box without reducing the font size. However, the cell height is not increased if for a Textline or Textflow the suboption *fitmethod=auto* is set, or if a Textflow is continued in another cell with the *continuetextflow* option.

This process of increasing the cell height applies only to cells containing a Textline or Textflow, but not for other types of cell contents, i.e. images, graphics, PDI pages, path objects, annotations, and fields.

You can use the *rowscalegroup* option to make sure that all rows in the same row scaling group are scaled to equal heights.

Continue the table in the next fitbox. If the table's resulting total height is larger than the fitbox (i.e. not all table cells fit into the fitbox), the formatter stops placing rows in the fitbox before it encounters the first row which doesn't fit into the fitbox.

If a cell spans multiple rows and not all of these fit into the fitbox, this cell is split. If the cell contains an image, PDI page, SVG graphics, path object, annotation, form field or Textline, the cell contents are repeated in the next fitbox unless *repeatcontent=false* has been specified. Textflows, however, are continued in the subsequent rows spanned by the cell (see figure 9.38).

You can use the *rowjoininggroup* option to make sure that all rows in a joining group appear together in a fitbox. All rows which belong to the header or footer plus one body row automatically form a row joining group. The formatter may therefore stop placing table rows before it encounters the first row which doesn't fit into the fitbox (see figure 9.37).

You can use the *return* option to make sure that no more rows are placed in the table instance after placing the affected row.

Split a row. A row may be split if it is very tall or if there is only a single body row. If the last body row doesn't fully fit into the fitbox, it is moved to the next fitbox. This behavior is controlled by the *minrowheight* option of *PDF_fit_table()* which has a default

value of 100%. If you reduce the *minrowheight* value the specified percentage of the content of the last body row is placed in the current fitbox and the rest of the row in the next fitbox (see figure 9.39).

You can check whether a row has been split with the *rowsplit* keyword of *PDF_info_table()*.

Shrink the table to the width of the fitbox. The table width as sum of the column widths may be larger than the fitbox width, e.g. after fitting a Textline horizontally. In this case all columns are evenly reduced until the table width equals the width of the fitbox. This shrinking process is limited by the *horshrinklimit* option. In order to avoid any horizontal shrinking use *horshrinklimit=100%*.

You can query the horizontal shrinking factor with the *horshrinking* keyword of *PDF_info_table()*.

If the *horshrinklimit* threshold is exceeded, i.e. the table would have to be squeezed too much horizontally, the following error message appears:

```
Table width $1 exceeds fitbox width $2 (required shrinking factor $3 is smaller than 'horshrinklimit')
```

Here *\$1* designates the calculated table width, *\$2* the width of the fitbox as supplied to *PDF_fit_table()*, and *\$3* the calculated shrinking percentage. If you run into this error you must supply a larger value for the width of the fitbox.

Shrink the table to the height of the fitbox. The table height for the header and footer rows plus at least one body row or row join group may be larger than the fitbox height. In this case all rows are evenly reduced until the table height equals the height of the fitbox. This shrinking process is limited by the *vertshrinklimit* option. In order to avoid any vertical shrinking use *vertshrinklimit=100%*.

You can query the vertical shrinking factor with the *vertshrinking* keyword of *PDF_info_table()*.

If the *vertshrinklimit* threshold is exceeded, i.e. the table would have to be squeezed too much vertically, the following error message appears:

```
Table height $1 exceeds fitbox height $2 (required shrinking factor $3 is smaller than 'vertshrinklimit')
```

Here *\$1* designates the calculated table height, *\$2* the height of the fitbox as supplied to *PDF_fit_table()*, and *\$3* the calculated shrinking percentage. If you run into this error you must supply a larger value for the height of the fitbox.

Subsequent fitboxes with different widths. If subsequent fitboxes are wider than the first fitbox, the table cells are not expanded to completely fill the fitbox width. As a result all table instances whose fitboxes are wider than the first fitbox have the same width by default. To completely fill the fitbox the option *fitmethod=auto* must be supplied. With *fitmethod=meet* the cell heights are suitably scaled to preserve their aspect ratio.

If subsequent fitboxes are narrower than the first fitbox, the widths of the table cells are reduced to fit into the narrower fitbox. In order to calculate the table width anew, call *PDF_fit_table()* with *rewind=1*.

9.4 Matchboxes

Matchboxes provide access to coordinates calculated by PDFlib as a result of placing some content on the page. Matchboxes are not defined with a dedicated API method, but with the *matchbox* option in the method call which places the actual element, for example *PDF_fit_textline()* and *PDF_fit_image()*. Matchboxes can be used for various purposes:

- ▶ Matchboxes can be decorated, e.g. filled with color or surrounded by a frame.
- ▶ Matchboxes can be used to automatically create one or more annotations with *PDF_create_annotation()*.
- ▶ Matchboxes define the height of a text line which is fit into a box with *PDF_fit_textline()* or the height of a text fragment in a Textflow which is decorated (*boxheight* option).
- ▶ Matchboxes define the clipping area for an image.
- ▶ The coordinates of the matchbox and other properties can be queried with *PDF_info_matchbox()* to perform some other task, e.g. insert an image.

For each element PDFlib calculates the matchbox as a rectangle corresponding to the bounding box which describes the position of the element on the page (as specified by all relevant options). For Textflows and table cells a matchbox may consist of multiple rectangles because of line or row breaking.

The rectangle(s) of a matchbox are drawn before drawing the element to be placed. As a result, the element may obscure the effect of the matchbox border or filling, but not vice versa. In particular, those parts of the matchbox which overlap the area covered by an image are hidden by the image. If the image is placed with *fitmethod=slice* or *fitmethod=clip* the matchbox borders outside the image fitbox will be clipped as well. To avoid this effect the matchbox rectangle can be drawn using the basic drawing functions, e.g. *PDF_rect()*, after the *PDF_fit_image()* call. The coordinates of the matchbox rectangle can be retrieved using *PDF_info_matchbox()* as far as the matchbox has been provided with a name in the *PDF_fit_image()* call.

In the following sections some examples for using matchboxes are shown. For details about the functions which support the *matchbox* option list, see the *PDFlib API Reference*.

Note Matchboxes are not supported in blind mode, i.e. formatting with the blind option.

9.4.1 Decorating a Textline

Let's start with a discussion of matchboxes in text lines. In *PDF_fit_textline()* the matchbox is the text box of the supplied text. The width of the text box is the text width, and the height is the capheight of the given font size, by default. To illustrate the matchbox size the following code fragment will fill the matchbox with blue background color (see figure 9.40a).

```
String optlist =  
    "font=" + normalfont + " fontsize=8 position={left top} " +  
    "matchbox={fillcolor={rgb 0.8 0.8 0.87} boxheight={capheight none}}";
```

```
p.fit_textline("Giant Wing Paper Plane", 2, 20, optlist);
```

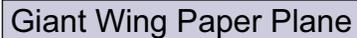
You can omit the *boxheight* option since *boxheight={capheight none}* is the default setting. It will look better if we increase the box height so that it also covers the descenders using the *boxheight* option (see figure 9.40b).

To increase the box height to match the font size we can use *boxheight={fontsize descender}* (see figure 9.40c).

In the next step we extend the matchbox by some offsets to the left, right and bottom to make the distance between text and box margins the same. In addition, we draw a rectangle around the matchbox by specifying the border width (see figure 9.40d).

Cookbook A full code sample can be found in the *Cookbook* topic `textflow/text_on_color`.

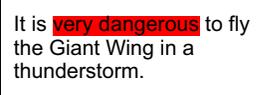
Fig. 9.40 Decorating a text line using a matchbox with various suboptions

Generated output	Suboptions of the matchbox option of <i>PDF_fit_textline()</i>
a) 	<code>boxheight={capheight none}</code>
b) 	<code>boxheight={ascender descender}</code>
c) 	<code>boxheight={fontsize descender}</code>
d) 	<code>boxheight={fontsize descender} borderwidth=0.3 offsetleft=-2 offsetright=2 offsetbottom=-2</code>

9.4.2 Using Matchboxes in a Textflow

Decorating parts of a Textflow. In this section we will decorate some text within a Textflow: The words *very dangerous* will be emphasized similar to a marker pen. To accomplish this the words are enclosed in the *matchbox* and *matchbox=end* inline options (see figure 9.41).

Fig. 9.41 Textflow with matchbox inline option

Generated output	Text and inline options for <i>PDF_create_textflow()</i>
	<code>It is <matchbox={fillcolor=red boxheight={ascender descender}}>very dangerous <matchbox=end> to fly the Giant Wing in a thunderstorm.</code>

Adding a Web link to the Textflow matchbox. Now we add a Web link to parts of a Textflow. In the first step we create the Textflow with a matchbox called *kraxi* indicating the text part to be linked. Second, we create the action for opening a URL. Third, we create an annotation of type *Link* with an invisible frame. In its option list we reference the *kraxi* matchbox to be used as the link's rectangle (the rectangle coordinates in *PDF_create_annotation()* are ignored if a matchbox is specified).

Cookbook A full code sample can be found in the *Cookbook* topic `textflow/weblink_in_text`.

```
/* create and fit Textflow with matchbox "kraxi" */
String tftext =
    "For more information about the Giant Wing Paper Plane see the Web site of " +
    "<underline=true matchbox={name=kraxi boxheight={fontsize descender}}>" +
    "Kraxi Systems, Inc.<matchbox=end underline=false>";
```

```
String optlist = "font=" + normalfont + " fontsize=8 leading=110%";
tflow = p.create_textflow(tftext, optlist);
if (tflow == -1)
    throw new Exception("Error: " + p.get_errmsg());

result = p.fit_textflow(tflow, 0, 0, 50, 70, "fitmethod=auto");
if (!result.equals("_stop"))
    { /* ... */ }

/* create URI action */
optlist = "url={http://www.kraxi.com}";
act = p.create_action("URI", optlist);

/* create Link annotation on matchbox "kraxi" */
optlist = "action={activate " + act + "} linewidth=0 usematchbox={kraxi}";
p.create_annotation(0, 0, 0, 0, "Link", optlist);
```

If the text *Kraxi Systems, Inc.* spans several lines the appropriate number of active rectangles for the link annotation is created with a single call to `PDF_create_annotation()` (see figure 9.42)

For information about
Giant Wing Paper
Planes see the Web
site of [Kraxi Systems
Inc.](http://www.kraxi.com)

<http://www.kraxi.com>

Fig. 9.42
Add Web links to a Textflow

9.4.3 Matchboxes and Images

Adding a Web link to an image. To add a Web link to the area covered by an image the image matchbox can be used. The code is similar to »Adding a Web link to the Textflow matchbox«, page 242, above. However, instead of placing the Textflow, fit the image using the following option list:

```
String optlist = "boxsize={130 130} fitmethod=meet matchbox={name=kraxi}";
p.fit_image(image, 10, 10, optlist);
```

Cookbook A full code sample can be found in the *Cookbook* topic `interactive/link_annotations`.

Drawing a frame around an image. In this example we want to use the image matchbox to draw a frame around the image. We completely fit the image into the supplied box while maintaining its proportions using `fitmethod=meet`. We use the `matchbox` option with the `borderwidth` suboption to draw a thick border around the image. The `strokecolor` suboption determines the border color, and the `linecap` and `linejoin` suboptions are used to round the corners.

Cookbook A full code sample can be found in the *Cookbook* topic `images/frame_around_image`.

The matchbox is always drawn before the image which means it would be partially hidden by the image. To avoid this we use the `offset` suboptions with 50 percent of the border width to enlarge the frame beyond the area covered by the image. Alternatively, we

could increase the border width accordingly. figure 9.43 shows the option list used with `PDF_fit_image()` to draw the frame.

Fig. 9.43 Using the image matchbox to draw a frame around the image

Generated output	Option list for <code>PDF_fit_image()</code>
	<pre> boxsize={60 60} position={center} fitmethod=meet matchbox={name=kraxi borderwidth=4 offsetleft=-2 offsetright=2 offsetbottom=-2 offsettop=2 linecap=round linejoin=round strokecolor={rgb 0.0 0.3 0.3}} </pre>

Align text at an image. The following code fragment shows how to align vertical text at the right margin of an image. The image is proportionally fit into the supplied box with a fit method of *meet*. The actual coordinates of the fitbox are retrieved with `PDF_info_matchbox()` and a vertical text line is placed relative to the lower right (x_2, y_2) corner of the fitbox. The border of the matchbox is stroked (see figure 9.44).

Cookbook A full code sample can be found in the *Cookbook* topic `images/align_text_at_image`.

```

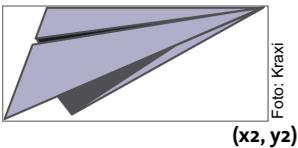
/* use this option list to load and fit the image */
String optlist = "boxsize={300 200} position={center} fitmethod=meet " +
    "matchbox={name=giantwing borderwidth=3 strokecolor={rgb 0.85 0.83 0.85}}";

/* load and fit the image */
...

/* retrieve the coordinates of the lower right (second) matchbox corner */
if ((int) p.info_matchbox("giantwing", 1, "exists") == 1)
{
    x2 = p.info_matchbox("giantwing", 1, "x2");
    y2 = p.info_matchbox("giantwing", 1, "y2");
}
/* start the text line at that corner with a small distance of 2 */
p.fit_textline("Foto: Kraxi", x2+2, y2+2, "font=" + font + " fontsize=8 orientate=west");

```

Fig. 9.44 Use the coordinates of the image matchbox to fit a text line

Generated output	Generation steps
	<p>Step 1: Fit image with matchbox</p> <p>Step 2: Retrieve matchbox info for coordinates (x_2, y_2)</p> <p>Step 3: Fit text line starting at retrieved coordinates (x_2, y_2) with option <code>orientate=west</code></p>

10 Interactive Features

Cookbook Code samples for creating interactive elements can be found in the interactive category of the PDFlib Cookbook.

10.1 Links, Bookmarks, and Annotations

This section explains how to create interactive elements such as bookmarks, form fields, and annotations. We will create a document with the following interactive elements:

- ▶ At the top right there is an invisible Web link to *www.kraxi.com* at the text *www.kraxi.com*. Clicking this area will bring up the corresponding Web page.
- ▶ A gray form field of type text is located below the Web link. Using JavaScript code it will automatically be filled with the current date.
- ▶ The red pushpin contains an annotation with an attachment. Clicking it will open the attached file.
- ▶ At the bottom left there is a form field of type button with a printer symbol. Clicking this button will execute Acrobat's menu item *File, Print*.
- ▶ The navigation page contains the bookmark »Our Paper Planes Catalog«. Clicking this bookmark will bring up a page of another PDF document.

In the next paragraphs we will show how to create these interactive elements.

Web link. Let's start with a link to the Web site *www.kraxi.com*. This is accomplished in three steps. First, we fit the text on which the Web link should work. Using the *matchbox* option with *name=kraxi* we specify the rectangle of the text's fitbox for further reference.

Second, we create an action of type *URI* (in Acrobat: *Open a web link*). This will provide us with an action handle which subsequently can be assigned to one or more interactive elements.

Third, we create the actual link with an annotation of type *Link*. The *action* option for the link contains the event name *activate* which will trigger the action, plus the *act* handle created above for the action itself. By default the link will be displayed with a thin black border. Initially this is convenient for precise positioning, but we disabled the border with *linewidth=0*.

```
normalfont = p.load_font("Helvetica", "unicode", "");
p.begin_page_ext(pagewidth, pageheight, "topdown");

/* place the text line "Kraxi Systems, Inc." using a matchbox */
String optlist =
    "font=" + normalfont + " fontsize=8 position={left top} " +
    "matchbox={name=kraxi} fillcolor={rgb 0 0 1} underline";
p.fit_textline("Kraxi Systems, Inc.", 2, 20, optlist);

/* create URI action */
optlist = "url={http://www.kraxi.com}";
int act = p.create_action("URI", optlist);

/* create Link annotation on matchbox "kraxi" */
optlist = "action={activate " + act + "} linewidth=0 usematchbox={kraxi}";
```

```

/* 0 rectangle coordinates will be replaced with matchbox coordinates */
p.create_annotation(0, 0, 0, 0, "Link", optlist);

p.end_page_ext("");

```

For an example of creating a Web link on an image or on parts of a Textflow, see Section 9.4, »Matchboxes«, page 241.

Cookbook A full code sample can be found in the *Cookbook* topic `interactive/link_annotations`.

Bookmark for jumping to another file. Now let's create the bookmark »Our Paper Planes Catalog« which jumps to another PDF file called *paper_planes_catalog.pdf*. First we create an action of Type *GoToR*. In the option list for this action we define the name of the target document with the *filename* option; the *destination* option specifies a certain part of the page which will be enlarged. More precisely, the document will be displayed on the second page (*page=2*) with a fixed view (*type=fixed*), where the middle of the page is visible (*left=50 top=200*) and the zoom factor is 200% (*zoom=2*):

```

String optlist =
    "filename=paper_planes_catalog.pdf " +
    "destination={page=2 type=fixed left=50 top=200 zoom=2}";

goto_action = p.create_action("GoToR", optlist);

```

In the next step we create the actual bookmark. The *action* option for the bookmark contains the *activate* event which will trigger the action, plus the *goto_action* handle created above for the desired action. The option *fontstyle=bold* specifies bold text and *textcolor=blue* makes the bookmark blue. The bookmark text »Our Paper Planes Catalog« is provided as a method parameter:

```

String optlist =
    "action={activate " + goto_action + "} fontstyle=bold textcolor=blue";
catalog_bookmark = p.create_bookmark("Our Paper Planes Catalog", optlist);

```

Clicking the bookmark displays the specified part of the page in the document.

Cookbook A full code sample can be found in the *Cookbook* topic `interactive/nested_bookmarks`.

Annotation with file attachment. In the next example we create a file attachment. We start by creating an annotation of type *FileAttachment*. The *filename* option specifies the name of the attachment, the option *mimetype=image/png* specifies its type. The annotation will be displayed as a pushpin (*iconname=pushpin*) in red (*annotcolor=red*) and has a tooltip (*contents={Get the Kraxi Paper Plane!}*). It will not be printed (*display=noprint*):

```

String optlist =
    "filename=kraxi_logo.png mimetype=image/png iconname=pushpin " +
    "annotcolor=red contents={Get the Kraxi Paper Plane!} display=noprint";
p.create_annotation(left_x, left_y, right_x, right_y, "FileAttachment", optlist);

```

The size of the symbol defined with *iconname* does not vary; the icon is displayed in its standard size in the top left corner of the specified rectangle.

10.2 Form Fields and JavaScript

Button form field for printing. The following example creates a button form field which can be used for printing the document. In the first version we add a caption to the button; later we use a printer symbol instead of the caption. We start by creating an action of type *Named* (in Acrobat: *Execute a menu item*). Also, we must specify the font for the caption:

```
print_action = p.create_action("Named", "menuname=Print");
button_font = p.load_font("Helvetica-Bold", "unicode", "");
```

The *action* option for the button form field contains the *up* event (in Acrobat: *Mouse Up*) as a trigger for executing the action, plus the *print_action* handle created above for the action itself. The *backgroundcolor=yellow* option specifies yellow background, while *bordercolor=black* specifies black border. The option *caption=Print* adds the text *Print* to the button, and *tooltip={Print the document}* creates an additional explanation for the user. The *font* option specifies the font using the *button_font* handle created above. By default, the size of the caption will be adjusted so that it completely fits into the button's area. Finally, the actual button form field is created with proper coordinates, the name *print_button*, the type *pushbutton* and the appropriate options:

```
String optlist =
    "action {up " + print_action + "} backgroundcolor=yellow " +
    "bordercolor=black caption=Print tooltip={Print the document} font=" +
    button_font;

p.create_field(left_x, left_y, right_x, right_y, "print_button", "pushbutton", optlist);
```

Now we extend the first version of the button by replacing the text *Print* with a little printer icon. To achieve this we load the corresponding image file *print_icon.jpg* as a template before creating the page. Using the *icon* option we assign the template handle *print_icon* to the button field, and create the form field similarly to the code above:

```
print_icon = p.load_image("auto", "print_icon.jpg", "template");
if (print_icon == -1)
{
    /* Error handling */
    return;
}
p.begin_page_ext(pagewidth, pageheight, "");
...
String optlist = "action={up " + print_action + "} icon=" + print_icon +
    " tooltip={Print the document} font=" + button_font;

p.create_field(left_x, left_y, right_x, right_y, "print_button", "pushbutton", optlist);
```

Cookbook A full code sample can be found in the *Cookbook* topic `form_fields/form_pushbutton`.

Simple text field. Now we create a text field near the upper right corner of the page. The user will be able to enter the current date in this field. We acquire a font handle and create a form field of type *textfield* which is called *date*, and has a gray background:

```

textfield_font = p.load_font("Helvetica-Bold", "unicode", "");
String optlist = "backgroundcolor={gray 0.8} font=" + textfield_font;
p.create_field(left_x, left_y, right_x, right_y, "date", "textfield", optlist);

```

By default the font size is *auto*, which means that initially the field height is used as the font size. When the input reaches the end of the field the font size is decreased so that the text always fits into the field.

Cookbook Full code samples can be found in the *Cookbook* topic `form_fields/form_textfield_height`.

Text field with JavaScript. In order to improve the text form field created above we automatically fill it with the current date when the page is opened. First we create an action of type *JavaScript* (in Acrobat: *Run a JavaScript*). The *script* option in the action's option list defines a JavaScript snippet which displays the current date in the *date* text field in the format month-day-year:

```

String optlist =
    "script={var d = util.printd('mmm dd yyyy', new Date()); " +
    "var date = this.getField('date'); date.value = d;}";

show_date = p.create_action("JavaScript", optlist);

```

In the second step we create the page. In the option list we supply the *action* option which attaches the *show_date* action created above to the trigger event *open* (in Acrobat: *Page Open*):

```

String optlist = "action={open " + show_date + "}";
p.begin_page_ext(pagewidth, pageheight, optlist);

```

Finally we create the text field as we did above. It will automatically be filled with the current date whenever the page is opened:

```

textfield_font = p.load_font("Helvetica-Bold", "winansi", "");
String optlist = "backgroundcolor={gray 0.8} font=" + textfield_font;
p.create_field(left_x, left_y, right_x, right_y, "date", "textfield", optlist);

```

Cookbook A full code sample can be found in the *Cookbook* topic `form_fields/form_textfield_fill_with_js`.

Formatting Options for Text Fields. In Acrobat it is possible to specify various options for formatting the contents of a text field, such as monetary amounts, dates, or percentages. This is implemented via custom JavaScript code used by Acrobat. PDFlib does not directly support these formatting features since they are not specified in the PDF reference. However, for the benefit of PDFlib users we present some information below which will allow you to realize formatting options for text fields by supplying simple JavaScript code fragments with the *action* option of *PDF_create_field()*.

In order to apply formatting to a text field JavaScript snippets are attached to the field as *keystroke* and *format* actions. The JavaScript code calls some internal Acrobat method where the parameters control details of the formatting.

The following sample creates two *keystroke* and *format* actions, and attaches them to a form field so that the field contents will be formatted with two decimal places and the EUR currency identifier:

```

keystroke_action = p.create_action("JavaScript",
    "script={AFNumber_Keystroke(2, 0, 3, 0, \"EUR \", true); }");

format_action = p.create_action("JavaScript",
    "script={AFNumber_Format(2, 0, 0, 0, \"EUR \", true); }");

String optlist = "font=" + font + " action={keystroke " + keystroke_action +
    " format=" + format_action + "}";
p.create_field(50, 500, 250, 600, "price", "textfield", optlist);

```

Cookbook A full code sample can be found in the *Cookbook* topic `form_fields/form_textfield_input_format`.

Validating form field input. The following sample attaches JavaScript to a form field as validate action to check whether the user input for a text field matches the required format *mm/dd/yyyy*:

```

optlist =
"script={" +
    "// JavaScript code for date mask format MM/DD/YYYY\n" +
    "var re = /^[0-9]{2}\\/[0-9]{2}\\/[0-9]{4}$/\n" +
    "if (event.value !=\"\") {\n" +
    "  if (re.test(event.value) == false) {\n" +
    "    app.alert ({\n" +
    "      cTitle: \"Incorrect Format\", \n" +
    "      cMsg: \"Please enter date using mm/dd/yyyy format\"\n" +
    "    });\n" +
    "  }\n" +
    "}\n" +
    "}\n" +
    "};";
validate_action = p.create_action("JavaScript", optlist);
textfield_font = p.load_font("Helvetica", "unicode", "");
optlist = "action={validate=" + validate_action + "} " +
    "backgroundcolor={gray 0.8} font=" + textfield_font;
p.create_field(11x, 11y, 11x, 11y, "startdate", "textfield", optlist);

```

10.3 Multimedia with Screen Annotations and Rendition Actions

General. Multimedia can be played in PDF documents with the help of and screen annotations and rendition actions. The media data can be embedded in the PDF document, taken from a separate external file, or pulled from a remote server. The media (called rendition) can be played at a fixed location on the page defined by an annotation, in a separate floating window, or in full screen mode. Optionally a controller can be displayed with the media so that the user can play, pause, resume or stop media playback. Alternatively these controls can be implemented with PDF form fields.

Playing PDF with multimedia elements requires Adobe Acrobat or a third-party PDF viewer with multimedia support. Most PDF viewers integrated in Web browsers currently don't support multimedia rendering.

The typical order of events for creating sound or movie with PDFlib is as follows:

- ▶ Load the video or sound file or URL with `PDF_load_asset()` and `type=Rendition`.
- ▶ Create a rendition action with `PDF_create_action()` and `type=Rendition`.

- ▶ Create a screen annotation which prepares the area on the page where the video will be played. The rendition action can be attached to the annotation for various trigger events, e.g. when the page is opened or when the annotation is clicked. Alternatively, you can attach the action to a bookmark or form field for controlling media playback.

The following code fragment implements the standard sequence with rendition, action and annotation:

```
// Create a rendition object with the video
rend_video = p.load_asset("Rendition", videofile,
    "mimetype=video/mp4 name=rendition_video controller");
if (rend_video == -1)
    throw new Exception("Error: " + p.get_errmsg());

// Create a rendition action for playing the video
act = p.create_action("Rendition",
    "target=annot_video operation=playfrombeginning rendition=" + rend_video);

// Create a screen annotation where the video will be played
p.create_annotation(25, 300, 575, 610, "Screen",
    "title=annot_video name=annot_video action={activate=" + act + "}");
```

As an alternative to this standard sequence you can play the rendition with JavaScript without any annotation (see »JavaScript for Multimedia«, page 251).

Cookbook A full code sample can be found in the *Cookbook* topic `multimedia/screen_annotations`.

Security considerations. Since multimedia playback has often been the target of attacks, Acrobat DC implements various security checks. In order to successfully play multimedia contents you must use one of the following methods to ensure that the PDF documents are trusted:

- ▶ Upon opening the document Acrobat displays a yellow bar at the top saying »Some features have been disabled to avoid potential security risks. Only enable these features if you trust this document«. In the »Options« menu click »Trust this document always«.
- ▶ Go to *Edit, Preferences, General..., Security (Enhanced)* and add the file or folder to the list of privileged locations.
- ▶ Apply a certification signature to the documents and enable *Edit, Preferences, General..., Security (Enhanced), Automatically trust documents with valid certification*.

MIME type. The `mimetype` option of `PDF_load_asset()` is required. It is crucial to supply the appropriate MIME type (content type) of the media to ensure reliable playback. The list of supported MIME types depends on the PDF viewer and version and the available media playback capabilities provided by the operating system. You can use the following JavaScript in Acrobat to determine the list of supported MIME types:

```
// Query all players and media types; use in JavaScript console
var mp = app.media.getPlayers("");
for (var i = 0; i < mp.length; i++) {
    console.println("\nMIME types supported by '" + mp[i].name + "':");
    console.println(mp[i].mimeTypes);
}
```

Table 10.1 lists MIME types for common sound and video formats. The *rendition* option of *PDF_create_action()* accepts a list of renditions where the PDF viewer picks the first rendition in a supported format. This feature can be used to supply a video in multiple formats to ensure that it can be played in different target environments.

Table 10.1 Common MIME types

Category	Common MIME types
Sound	audio/wav, audio/midi, audio/mp3, audio/mp4
Video	video/mpeg, video/msvideo, video/avi, video/mp4, video/quicktime

JavaScript for Multimedia. The JavaScript API for Acrobat offers ample opportunities for working with multimedia content in PDF. For example, instead of activating a video with the *operation* option of *PDF_create_action()* you can supply a simple JavaScript statement as an alternative to the *operation* option:

```
act = p.create_action("Rendition",
    "target=my_annot script={app.media.openPlayer();} rendition=" + rend_video);
```

JavaScript can also be used to manipulate renditions provided these are stored in the document's »rendition name tree« (in Acrobat such renditions are called »accessible to JavaScript«). This can be achieved with the *nametree* option of *PDF_load_asset()*. Subsequently the rendition doesn't require any annotation, but can be played with a JavaScript action, e.g. triggered by clicking a bookmark:

```
// Retrieve rendition by name and play it in a floating window with specific properties
// This works only for renditions in the name tree
```

```
var rendition = this.media.getRendition("rendition_sound");

if (rendition == null)
    app.alert("Couldn't find rendition");
else
{
    var floating = {
        width: 400,
        height: 300
    };
    var player = app.media.openPlayer({
        rendition: rendition,
        settings: {
            windowType: app.media.windowType.floating,
            floating: floating
        }
    });
}
```

The following code assumes that the JavaScript code above is stored in the String variable *js*, loads it in a virtual (PVF) file, creates a JavaScript asset and then a JavaScript action which is attached to a bookmark:

```
// Create a hidden rendition with the sound file and store it in the name tree
// to make it accessible to JavaScript
int rend_sound = p.load_asset("Rendition", soundfile,
```

```

        "mimetype=audio/midi style=hidden nametree name=rendition_sound");
if (rend_sound == -1)
    throw new Exception("Error: " + p.get_errmsg());

// Load JavaScript string into a virtual (PVF) file
p.create_pvf("/pvf/javascript/script1.js", js.getBytes(), "");

// Create a JavaScript asset from PVF file
int js_asset = p.load_asset("JavaScript", "/pvf/javascript/script1.js", "");
if (js_asset == -1)
    throw new Exception("Error: " + p.get_errmsg());

// Create a JavaScript action from the asset
int act_js = p.create_action("JavaScript", "javascript=" + js_asset);

// Bookmark for playing the sound rendition
p.create_bookmark("play sound with JavaScript (click bookmark)",
    "action={activate=" + act_js + "}");

```

By attaching two actions to a bookmark you can even play sound and video renditions at the same time, even without any corresponding Screen annotation. Renditions without any annotation are especially useful for playing sound which doesn't need an area on the page.

Cookbook A full code sample for sound and video renditions without annotations can be found in the *Cookbook* topic `multimedia/javascript_for_renditions`. The topic `interactive/renditions` in the *pCOS Cookbook* demonstrates how to enumerate all renditions in the name tree.

10.4 Geospatial PDF

Cookbook A full code sample can be found in the *Cookbook* topic `geospatial/starter_geospatial`.

10.4.1 Using georeferenced PDF in Acrobat

PDF 1.7ext3 allows geospatial reference information (world coordinates) to be added to PDF page contents. Geospatially referenced PDF documents can be used in Acrobat for several purposes. In Acrobat and Adobe Reader DC you must enable *Tools, Measure*:

- ▶ Display the coordinates of the map point under the mouse cursor: *Geospatial Location Tool*. You can copy the coordinates of the map point under the mouse cursor by right-clicking and selecting *Copy Coordinates to Clipboard*;
- ▶ Search for a location on the map: *Geospatial Location Tool*, right-click and select *Find a Location*, and enter the desired coordinates;
- ▶ Mark a location on the map: *Geospatial Location Tool*, right-click and select *Mark Location*;
- ▶ Measure distance, perimeter and area on geographic maps: *Measuring Tool*;

Only the first two functions mentioned above are available in Adobe Reader. Various settings for geospatial measuring can be changed in *Edit, Preferences, [General...], Measuring (Geo)*, e.g. the preferred coordinate system for coordinate readouts.

Geospatial features in PDFlib are implemented with the following functions and options:

- ▶ One or more georeferenced areas can be assigned to a page with the *viewports* option (and suboption *georeference*) of *PDF_begin/end_page_ext()*. Viewports allow different

geospatial references (specified by the *georeference* option) to be used on different areas of the page, e.g. for multiple maps on the same page. This method works in all PDF viewers with support for georeferenced PDF.

Cookbook A full code sample can be found in the *Cookbook* topic `geospatial/starter_geospatial`.

- ▶ The *georeference* option of `PDF_load_image()` can be used to assign an earth-based coordinate system to an image.

Cookbook A full code sample can be found in the *Cookbook* topic `geospatial/georeferenced_image`.

- ▶ The *georeference* option of `PDF_open_pdi_page()`, `PDF_load_graphics()` and `PDF_begin_template_ext()` can be used to assign an earth-based coordinate system to a Form XObject. However, this method is not recommended since it is not supported in any known viewer including Acrobat DC.

10.4.2 Geographic and projected Coordinate Systems

A geographic coordinate system describes the earth in geographic coordinates, i.e. angular units of latitude and longitude. A projected coordinate system can be specified on top of a geographic coordinate system and describes the transformation of points in geographic coordinates to a two-dimensional (projected) coordinate system. The resulting coordinates are called Northing and Easting values; degrees are no longer required for projected coordinate systems. While geographic coordinate systems are in use for GPS and other global applications, projections are required for map-making and other applications with more or less local character.

For historical and mathematical reasons a variety of different coordinate systems is in use around the world. Both geographic and projected coordinate systems can be described using two well-established methods which are called EPSG and WKT. Although both WKT and EPSG can be used in Acrobat, Acrobat does not implement all possible EPSG codes. In particular, EPSG codes for geographic coordinate systems don't seem to be supported in Acrobat. In this case the use of WKT is recommended. The following Web site delivers the WKT corresponding to a particular EPSG code:

www.spatialreference.org/ref/epsg

10.4.3 Coordinate System Examples

Examples for geographic coordinate systems. The WGS84 (World Geodetic System) geographic coordinate system is the basis for GPS and many applications (e.g. *OpenStreetMap*). It can be expressed as follows in the *worldsystem* suboption of the *georeference* option:

```
worldsystem={type=geographic wkt={
GEOGCS["WGS 84",
  DATUM["WGS_1984", SPHEROID["WGS 84", 6378137, 298.257223563]],
  PRIMEM["Greenwich", 0],
  UNIT["degree", 0.01745329251994328]]
}}
```

The ETRS (European Terrestrial Reference System) geographic coordinate system is almost identical to WGS84. It can be specified as follows:

```

worldsystem={type=geographic wkt={
GEOGCS["ETRS_1989",
    DATUM["ETRS_1989", SPHEROID["GRS_1980", 6378137.0, 298.257222101]],
    PRIMEM["Greenwich", 0.0],
    UNIT["Degree", 0.0174532925199433]]
}}

```

Note EPSG codes for the WGS84 and ETRS systems are not shown here because Acrobat doesn't seem to support EPSG codes for geographic coordinate systems, but only for projected coordinate systems (see below).

Examples for projected coordinate systems. A projection is based on an underlying geographic coordinate system. In the following example we specify a projected coordinate system suitable for use with GPS coordinates.

In middle Europe the system called ETRS89 UTM zone 32 N applies. It uses the common UTM (Universal Mercator Projection), and can be expressed as follows in the *worldsystem* suboption of the *georeference* option:

```

worldsystem={type=projected wkt={
PROJCS["ETRS_1989_UTM_Zone_32N",
    GEOGCS["GCS_ETRS_1989",
        DATUM["D_ETRS_1989", SPHEROID["GRS_1980", 6378137.0, 298.257222101],
            TOWGS84[0, 0, 0, 0, 0, 0, 0]],
        PRIMEM["Greenwich", 0.0],
        UNIT["Degree", 0.0174532925199433]],
    PROJECTION["Transverse_Mercator"],
    PARAMETER["False_Easting", 500000.0],
    PARAMETER["False_Northing", 0.0],
    PARAMETER["Central_Meridian", 9.0],
    PARAMETER["Scale_Factor", 0.9996],
    PARAMETER["Latitude_Of_Origin", 0.0],
    UNIT["Meter", 1.0]]
}}

```

The corresponding EPSG code for this coordinate system is 25832. As an alternative to WKT, the system above can also be specified via its EPSG code as follows:

```

worldsystem={type=projected epsg=25832}

```

10.4.4 Georeferenced PDF Restrictions in Acrobat

We experienced the following shortcomings when working with georeferenced PDF in Acrobat DC:

- ▶ EPSG codes don't seem to work at all for geographic coordinate systems, but only for projected systems.
Workaround: use the corresponding WKT instead of the EPSG code.
- ▶ Attaching geospatial data to Form XObjects does not work. For this reason the *georeference* option for *PDF_open_pdi_page()*, *PDF_begin_template_ext()* and *PDF_load_graphics()* is not recommended although it should work according to the PDF Reference.
Workaround for creating vector-based maps: you can attach the geospatial data to the page, i.e. use the *viewports* option of *PDF_begin_page_ext()*.

- ▶ Overlapping maps: you can place multiple image-based maps on the same page. If the maps overlap and you display the coordinates of a point in the overlapping area, Acrobat uses the coordinates of the map which has been placed last (this makes sense since this is also the map which is visible). However, if both image handles are identical (i.e. retrieved with a single call to `PDF_load_image()`), Acrobat does no longer take into account the different image geometries: the coordinates of the first image are incorrectly extended to the area of second image, resulting in wrong coordinate readouts.
Workaround: if you need multiple instances of the same image-based map on the same page, open the image multiply.
- ▶ The area measurement tool doesn't work correctly for geographic coordinate systems, but only for projected systems.



11 Tagged PDF and PDF/UA

11.1 Tagged PDF Basics

Tagged PDF is a requirement for the PDF/UA, PDF/A-1a, PDF/A-2a, and PDF/A-3a ISO standards, Section 508 in the U.S., BITV in Germany and many other regulations. Tagged PDF enhances PDF with document structure information which offers the following advantages:

- ▶ **Accessibility:** Tagged PDF is accessible for users with disabilities, e.g. via Acrobat's built-in Read Aloud feature or more advanced screen reader software.
- ▶ **Reliable export and conversion to other document formats:** converting Tagged PDF to other formats such as RTF, XML or HTML results in more accurate output.

PDF/UA enhances Tagged PDF by specifying requirements for the document tags. If you want to create accessible PDF documents we recommend to follow the additional rules for PDF/UA; see Section 11.4, »PDF/UA for Universal Accessibility«, page 293, for details.

If you cannot obey all PDF/UA requirements (e.g. because you must assemble documents from existing PDFs which themselves don't conform to PDF/UA) we recommend to disable PDF/UA mode and obey as many PDF/UA rules as possible.

Cookbook Code samples for generating Tagged PDF can be found in the pdfua category of the PDFlib Cookbook. All Tagged PDF samples in the Cookbook create PDF/UA.

11.1.1 Creating the Logical Structure Tree (Structure Hierarchy)

Tagged PDF can only be created if the client provides information about the document's internal structure and obeys certain rules when generating PDF output. In order to create Tagged PDF the *tagged* document option must be set to *true*, and the *lang* option is recommended (required for PDF/UA-1):

```
if (p.begin_document("tagged.pdf", "tagged=true lang=en") == -1) {  
    throw new Exception("Error: " + p.get_errmsg());  
}
```

The logical structure in a Tagged PDF document is described by a hierarchy of elements, called the structure hierarchy or logical structure tree. Starting at the root level the structure hierarchy consists of an arbitrary number of levels. On each level an element may contain zero or more items of the following kinds:

- ▶ Other structure elements, e.g. the *Document* element may contain multiple *Art* (article) elements. Each *Art* element in turn may contain multiple *P* (paragraph) elements.
- ▶ Content items, i.e. sequences of text and graphics on the page, XObjects created from imported images, and OBJR references to annotations and form fields. These items represent the content associated with a structure element.

Structure hierarchy in Acrobat. You can display tag names and structure hierarchy in Acrobat DC as follows:

- ▶ select *View, Show/Hide, Navigation Panes, Tags* (see figure 11.1)

Tag nesting functions. Structure elements can be created explicitly with the *PDF_begin_item()* method which must always be paired with a corresponding call to

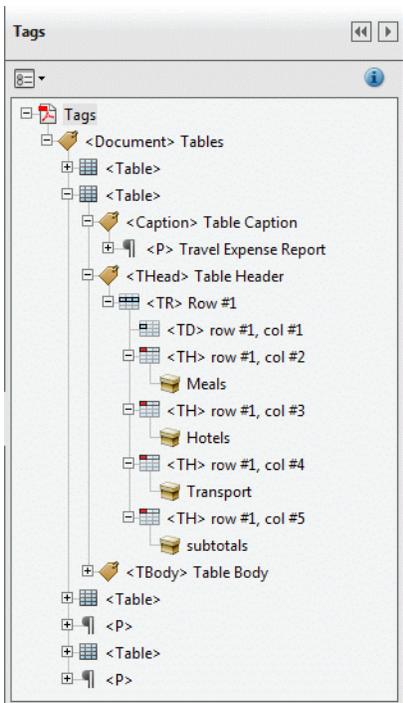


Fig. 11.1
Acrobat's tag pane displays the document's logical structure tree

`PDF_end_item()`. For example, the code fragment below creates a hierarchy consisting of a section which contains a heading and a paragraph. The hierarchical relationship is visualized with indentation:

```
/* Create a structure element of type "Sect" (section) */
id_sect = p.begin_item("Sect", "Title={Past and Future}");

    /* Create a structure element of type "H1" (heading) */
    id_h1 = p.begin_item("H1", "Title={Company History}");
        p.fit_textline(...);
    p.end_item(id_h1);

    /* Create a structure element of type "P" (paragraph) */
    id_p = p.begin_item("P", "");
        p.fit_textline(...);
    p.end_item(id_p);

/* Close "Sect" */
p.end_item(id_sect);
```

By default, structure elements are inserted as a child of the currently active item after all other child items which may already be present. This leads to the correct tree structure if the elements are created in logical order. See Section 11.2.4, »Creating Contents out of Order«, page 288, for more advanced techniques.

Structure elements may carry one or more attributes which are supplied via options, e.g. `lang` for specifying the natural language of a content element or `Alt` for alternative text for images. The set of available options depends on the structure type.

Abbreviated tagging. In many situations the content of a structure element can be created with a single call to a PDFlib fitting method, resulting in the typical sequence `PDF_begin_item()/PDF_fit_*/PDF_end_item()`. This sequence can be reduced with abbreviated tagging. Placing content and supplying tagging information can be combined in a single method call using the *tag* option which is supported by most functions for creating page content. The code fragment above can be simplified by supplying the *tag* option to the text and image placement functions:

```
/* Create a structure element of type "Sect" (section) */
id_sect = p.begin_item("Sect", "Title={Past and Future}");

    /* Create a structure element of type "H1" (heading) */
    p.fit_textline(..., "tag={tagname=H1 Title={Company History}}");

    /* Create a structure element of type "P" (paragraph) */
    p.fit_textline(..., "tag={tagname=P}");

/* Close "Sect" */
p.end_item(id_sect);
```

Since the tag nesting method `PDF_begin_item()` also supports the *tag* option both methods can be combined. This is useful in situations where a tag requires a nested tag which in turn contains multiple content items. For example, a *Caption* element containing a *P* element which in turn contains multiple text items can be created as follows:

```
id_caption = p.begin_item("Caption", "tag={tagname=P}");
    p.fit_textline(...);
    p.fit_textline(...);
p.end_item(id_caption);
```

A second *tag* option can also be nested as suboption in the option list of a *tag* option. Indeed, *tag* options can be nested to an arbitrary level. If a nested structure element contains only a single content item the code fragment above can be further simplified by applying abbreviated tagging directly in the call for placing the text on the page:

```
p.fit_textline(..., "tag={tagname=Caption tag={tagname=P}}");
```

Since abbreviated tagging does not expose the handle of the generated structure element it cannot be used with `PDF_activate_item()` or the *parent* suboption of the *tag* option since these require a structure element id. The detailed operation of the *tag* option is as follows:

- ▶ A new structure element is created for the generated content, and closed before returning from the call. The following situations are exceptions from this rule:
 - ▶ A structure element created with the *tag* option in `PDF_begin_document()` will be closed in `PDF_end_document()`.
 - ▶ When multiple tags are supplied with the *tag* option of `PDF_begin_item()` all of these tags will be closed in the corresponding call to `PDF_end_item()`.
- ▶ `PDF_fit_table()`: *tagname=Table* or a tag name which is role-mapped to *Table* in `PDF_fit_table()` instructs PDFlib to create the required table tags (see Section 11.2.1, »Automatic Table Tagging«, page 278). Automatically generated *TH* and *TD* tags for table cells can be further qualified with the *tag* option of `PDF_add_table_cell()`.

- ▶ `PDF_fit_textflow()`: the complete Textflow instance forms the new structure element. Abbreviated tagging has priority over tags in Textflow option lists which are discussed below.
- ▶ The generated element is a child of the currently active item or the item which has been supplied in the *parent* option.
- ▶ Grouping elements can only be created with `PDF_begin_item()`, but not with abbreviated tagging except `PDF_begin_document()`.
- ▶ Except for `PDF_begin_document()` and `PDF_add_table_cell()` abbreviated tagging can only be used in *page* scope.

Tags in Textflow option lists. As an alternative to explicit tagging with `PDF_begin/end_item()` or abbreviated tagging with *tag* option of many functions, tagging information can also be supplied as part of a Textflow. Structure element names and attributes can be supplied as options in `PDF_create_textflow()` and `PDF_add_textflow()`. Tagging information can be included directly in Textflow option lists:

```
<tagbegin={tagname=H1 Title={Company History}}>Introduction<tagend>\n
<tagbegin={tagname=P}>Kraxi Systems was founded in ... <tagend>\n
```

The *tagbegin* option accepts the same options as abbreviated tagging, while the *tagend* option closes the most recently opened structure element. These tagging options can be combined with other Textflow control options.

All instances of a Textflow with tags must be placed on output pages, i.e. `PDF_fit_textflow()` must be called until it returns *_stop*. In the rare situation that not all Textflow instances are placed, i.e. some part of the text is lost, the *tagstop* option must be supplied to the last call of `PDF_fit_textflow()` for this Textflow.

When placing a Textflow in the output, one or more Textflow instances are created depending on the amount of text and size of the fitbox. PDFlib ensures that all Textflow instances are children of a common structure element and that an element which is interrupted at the end of a Textflow instance is automatically activated again in the next instance. The placement of Textflow instances doesn't matter, i.e. on the same page or on different pages.

Tags in Textflow option lists are ignored in the following situations:

- ▶ in contexts where tagging is not supported, e.g. on templates or with untagged output;
- ▶ if the *tag* option has been supplied to `PDF_fit_textflow()`.

Textflow macros (see Section 9.2.1, »Inline Option Lists and Macros«, page 209) may be used to define convenient shortcuts for tagging, e.g.

```
macro {
  H1 {tagbegin={tagname=H1 Title = {Company History}} }
    /H1 {tagend}

  P {tagbegin={tagname=P} }
    /P {tagend}
}
```

The macros can then be used in inline option lists using the `&` introducer for macros:

```
<&H1>Introduction<&/H1>\n
<&P>Kraxi Systems was founded in ... <&/P>\n
```

Keep in mind that option lists always deliver only a single value of a particular option. If an option is used more than once, the last occurrence overrides all previous instances of that option in the same option list. For example, the following construct doesn't start two structure elements since the second *tagbegin* overrides the first; only the *LI* item is created while the *L* item is ignored:

```
<tagbegin={tagname=L} tagbegin={tagname=LI}> // wrong: single option list!
```

In order to create two structure elements you must supply the two *tagbegin* options in separate inline option lists or in separate calls to *PDF_add_textflow()*, e.g.

```
<tagbegin={tagname=L}><tagbegin={tagname=LI}> // correct: two option lists
```

If two nested structure elements enclose the same contents you can supply the *tagbegin* option with another nested *tag* option in a single option list. For example, since a *Caption* element cannot contain content items it must contain another element such as *P* which holds the actual contents. This construct can be expressed with inline Textflow options as follows:

```
<tagbegin={tagname=Caption tag={tagname=P}}>Figure 9.1<tagend>
```

The *tagend* option closes the corresponding *Caption* element as well as the nested *P* element.

Cookbook A code sample for creating Tagged PDF with tags in Textflow option lists can be found in the topic *inline_tags_in_textflow_pdfua1* topic in the *pdfua* category of the *PDFlib Cookbook*.

11.1.2 Standard Structure Types

PDF 1.7 and PDF 2.0 tagsets. PDF supports standard structure types which are designed to support a wide range of document classes. PDFlib supports all standard elements types of PDF 1.x and PDF 2.0 according to Table 11.1. The descriptions provided in the table are intended to assist in selecting appropriate types. Depending on the selected PDF output compatibility different groups of structure types, so-called tagsets, are available:

- ▶ The PDF 1.7 tagset is always available for *compatibility=1.7* and below. For *compatibility=2.0* this tagset is available if *pdftagset=1.7*. It includes all types which are not marked as *PDF 2.0* in Table 11.1.
- ▶ The PDF 2.0 tagset is default for *compatibility=2.0*. It includes most of the types in the PDF 1.7 tagset plus the types marked as *PDF 2.0* in Table 11.1.

All structure elements must either be drawn from the set of standard structure types, or (in case of custom elements) must be role mapped to one of the standard structure types (see Section 11.1.5, »Custom Structure Types, Role Maps and custom Tagsets«, page 271). The standard structure types fall into several categories including the following:

- ▶ Grouping level structure types are containers which group other elements, but don't contain any page content.
- ▶ Block level structure types describe the overall layout of content on the page.
- ▶ Inline level structure types describe the layout of content within a Block level element.
- ▶ Special structure types for lists, tables, Japanese Ruby/Warichu, and others.

- ▶ Pseudo structure types don't create any structure element, but are used for marking up content with certain characteristics.

Some structure types may be used within different categories.

Ultimately the page contents are described by content items, i.e. text, images, vector graphics and interactive elements. The relationship of structure types as well as the elements which may contain content items are governed by the tag nesting rules. Although there are some exceptions, the nesting rules for standard structure types can roughly be summarized as follows:

- ▶ Document level and Grouping level elements contain only other Document level, Grouping or Block level elements, but not any Inline level elements or content items.
- ▶ Content items can only be contained in Block level or Inline level elements, and the appropriate sub-elements of lists and tables.
- ▶ Block level elements may contain only Inline level elements and content items.
- ▶ Inline level elements may contain only other Inline level elements and content items.

The structure element nesting rules for PDF 1.7 and PDF 2.0 are detailed in the next sections. Applications must obey to these nesting rules. Additional rules apply to PDF/UA-1 (see Section 11.4, »PDF/UA for Universal Accessibility«, page 293).

The nesting rule checks for structure elements can be disabled or relaxed with the *checktags* option of `PDF_begin_document()`. However, this is not recommended since it may result in an invalid structure hierarchy. This option is intended as a migration aid for legacy applications. Some rules in Table 11.3 are marked as »strict rule«. The option *checktags=relaxed* enforces all checks except the strict rules.

Table 11.1 Standard types (tags) in Tagged PDF and pseudo types added by PDFlib

structure type	description
Document level structure types (can contain document-level, grouping or block-level elements, but not inline-level elements or content items)	
Document	Self-contained logical document. In PDF 2.0 required as root element of the structure tree. It may contain nested Document elements.
Document-Fragment	(PDF 2.0) Logical document fragment, typically created by extracting a part from an original complete document. It may contain nested DocumentFragment elements.
Grouping level structure types (cannot contain page content)	
Art	(Article; requires pdftagset=1.7 if compatibility=2.0) Article within a document.
Aside	(PDF 2.0) Content that is distinct from other content within its parent structure element, e.g. text in a sidebar or auxiliary comment.
BlockQuote	(Block quotation; requires pdftagset=1.7 if compatibility=2.0) Portion of text consisting of one or more paragraphs attributed to someone other than the author of the surrounding text
Div	(Division) Division of a document without semantic intent. It can also be used for non-semantic grouping, such as associating a lang attribute with a sequence of block-level elements, or for role-mapping a custom grouping structure element for which none of the other standard structure elements are suitable.
Index	(Index; requires pdftagset=1.7 if compatibility=2.0) Sequence of entries containing identifying text accompanied by Reference elements that point out occurrences of the specified text in the main body of the document

Table 11.1 Standard types (tags) in Tagged PDF and pseudo types added by PDFlib

structure type	description
NonStruct	(Non-structural element) Grouping element having no inherent structural significance; it serves solely for grouping purposes.
Part	Encloses a grouping of structure elements without consideration for their hierarchy (semantic equivalent of Div).
Private	(Private element; requires pdftagset=1.7 if compatibility=2.0) Grouping element containing private content belonging to the application
Sect	(Section) Encloses a grouping of structure elements with consideration for their hierarchy.
TOC	(Table of contents; requires pdftagset=1.7 if compatibility=2.0) List made up of table of contents items (structure type TOCI)
TOCI	(Table of contents item; requires pdftagset=1.7 if compatibility=2.0) Member of a table of contents. It should contain a suitable Link element.
Block level structure types (can contain only inline-level elements)	
FENote	(PDF 2.0; Grouping, Block or Inline level) Contains footnotes and endnotes which are not normally read as part of the enclosing content from which they are referenced.
H	(Heading; in PDF/UA-1 only with structuretype=strong) Low-level division of content usually referred to as a heading. H elements are intended for hierarchical nesting. Weak structuring with H1 etc. is preferable.
H1...H6	Low-level division of content usually referred to as a heading. These types should be used if the application cannot hierarchically nest sections and thus cannot determine the level of a heading from its level of nesting. The Lb1 structure type can be used to enclose section numbering within a heading.
H7, H8...	(PDF 2.0 or PDF/UA-1 with role mapping) Additional heading levels for deeper nesting. In PDF/UA-1 role mapping to P is recommended.
P	(Paragraph) Low-level division of content. Although in many cases it will be used for paragraphs it may enclose any low-level division of content.
Title	(PDF 2.0; Grouping or Block level) Encloses content that is usually referred to as the title of a document or high-level division of content. There should be a reference from the enclosing content to the footnote or endnote.
Inline level structure types	
BibEntry	(Bibliography entry); requires pdftagset=1.7 if compatibility=2.0) Reference identifying the external source of some cited content
Code	(Requires pdftagset=1.7 if compatibility=2.0) Fragment of computer program text
Em	(PDF 2.0) Encloses content for the purpose of emphasis. Nested Em elements may be used to increase the level of emphasis.
Lbl	Content that distinguishes it from other content inside the same parent element, e.g. bullets or numbers for a list item, numbers in a caption, enumeration in a heading.
Note	(Requires pdftagset=1.7 if compatibility=2.0) Item of explanatory text, such as a footnote or an endnote, that is referred to from within the body of the document. This element may have Lbl as a child.
Sub	(PDF 2.0) Encloses content typically perceived as a sub-division inside a Block level element, e.g. the verse in a poem, a line of source code, or a line in a postal address.
Quote	(Quotation; requires pdftagset=1.7 if compatibility=2.0) Inline portion of text attributed to someone other than the author of the surrounding text. The quoted text should be contained inline within a single paragraph.

Table 11.1 Standard types (tags) in Tagged PDF and pseudo types added by PDFlib

structure type	description
Reference	(Requires pdftagset=1.7 if compatibility=2.0) Citation to content elsewhere in the document. It should be used for local links.
Span	(Not related to row or column spans in tables) Generic inline portion of content having no particular inherent characteristics. It can be used to delimit a range of text with a certain attribute, e.g. lang.
Strong	(PDF 2.0) Encloses content for the purpose of strong importance, seriousness or urgency.
Structure types for interactive features (Grouping, Block or Inline level) (see Section 11.2.1, »Tagging Interactive Elements«, page 281)	
Annot	(PDF 1.5) Associates content with one or more annotations. It can also be used to include one or more annotations in the structure tree. This element must be used for annotations unless Link or Form is more appropriate.
Form	Associates content with a corresponding form field. It can also be used to include a form field in the structure tree. A Lbl element may be used to mark up a form field's label.
Link	Associates content with a Link annotation. It can also be used to include a Link annotation in the structure tree. This element must be used for Link annotations.
Structure types for Japanese Ruby and Warichu (Inline level)	
Ruby	(PDF 1.5) A side note written in a smaller text size and placed adjacent to the base text to which it refers. The Ruby element serves as wrapper around the entire ruby assembly. It must contain either an RB/RT pair or a sequence RP, RT, RP.
RB	(Ruby base text; PDF 1.5) Full-size text to which the ruby annotation is applied.
RT	(Ruby annotation text; PDF 1.5) Smaller-size text that shall be placed adjacent to the ruby base text.
RP	(Ruby punctuation) Punctuation surrounding the ruby annotation text. It is used only when a ruby annotation cannot be properly formatted in a ruby style and instead is formatted as a normal comment, or when it is formatted as a warichu.
Warichu	(Warichu; PDF 1.5) Comment or annotation in a smaller text size and formatted onto two smaller lines within the height of the containing text line and placed following (inline) the base text to which it refers. The Warichu element serves as wrapper around the entire warichu assembly. It must contain the sequence WP, WT, WP.
WT	(Warichu text; PDF 1.5) Smaller-size text of a warichu comment that is formatted into two lines and placed between surrounding WP elements
WP	(Warichu punctuation; PDF 1.5) Punctuation that surrounds the content in the WT element
List structure types	
L	(List; Block or Inline level) Sequence of items that are semantically related to each other
LI	(List item) Content for an individual member of a list
LBody	(List body) Actual content of a list item
Table structure types (tables can be tagged automatically, see Section 11.2.1, »Automatic Table Tagging«, page 278)	
Table	(Block or Inline level) Two-dimensional logical structure of rectangular cells, possibly including a complex substructure
TR	(Table row) Row of table header cells (TH) or table data cells (TD)
TH	(Table header cell) Table header cell describing one or more rows, columns or rows and columns of the table.
TD	(Table data cell) Table cell containing data that is part of the table's content

Table 11.1 Standard types (tags) in Tagged PDF and pseudo types added by PDFlib

structure type	description
THead	(Table header row group; PDF 1.5) Group of TR elements that constitute the header of a table
TBody	(Table body row group; PDF 1.5) Group of TR elements that constitute the main body portion of a table
TFoot	(Table footer row group; PDF 1.5) Group of TR elements that constitute the footer of a table
Caption structure type (Grouping or Block level)	
Caption	(PDF 2.0; in PDF 1.5 only for tables and lists) Brief portion of text describing a table, list, figure or other element
Figure and Formula structure types (Grouping, Block or Inline level)	
Figure	Graphical content
Formula	Mathematical, chemical or other formula. This type identifies an entire content element as a formula.
Artifact structure type (Grouping, Block or Inline level)	
Artifact	(PDF 2.0) Artifact, to be distinguished from real page content (see Section 11.1.6, »Artifacts«, page 272). Artifacts as elements of the structure tree can only be used in PDF 2.0; Artifacts as marked content can be used in any PDF version. Artifacts within the structure tree may be used for irrelevant content which requires other structure elements as context, e.g. line numbers. This type is also required to tag annotations as artifacts.
Pseudo structure types	
Artifact	Artifact as direct marked content, to be distinguished from real page content (see Section 11.1.6, »Artifacts«, page 272).
ASpan	(Accessibility span; written to PDF as Span, but must be distinguished from the inline level structure type Span) Attach accessibility properties to content which does not belong to a structure element or which resembles a part of a structure element. The ASpan pseudo element is written as Span with an accessibility attribute such as Alt, ActualText, Lang, or E. An ASpan is not associated with any structure element.
ReversedChars	Specifies text in a right-to-left script with reversed characters.
Clip	Specifies a marked clipping sequence. This is a sequence containing only clipping paths or text in text rendering mode 7, but no visible graphics or <code>PDF_save()</code> / <code>PDF_restore()</code> .

Regular elements vs. direct elements. Most structure elements are emitted to PDF by marking up the corresponding text or graphics and adding a corresponding entry in the document structure tree. Such elements are called regular elements and are displayed in the structure tree in Acrobat's tags pane.

In contrast, some elements may consist only of marked content without an entry in the structure tree. Such elements are not visible in Acrobat's tags pane. The advantage of such direct elements is that they require fewer space in the PDF output. The status of some structure types can be changed with the *direct* option (the default is *true*). Table 11.2 compares regular structure elements and direct elements. The table also indicates additional nesting rules for direct elements in PDF 2.0 (in both cases the usual PDF 2.0 nesting rules apply to the child elements).

The option *direct=false* is required in cases where an element which is direct by default contains another indirect element as child, e.g. *Reference* contains *Link*.

Empty structure elements. In general, structure elements which contain neither child elements nor content items should be avoided. However, there are some exceptions to

Table 11.2 Regular and direct structure elements

	<i>regular structure elements (direct=false)</i>	<i>direct structure elements (direct=true)</i>
<i>affected elements</i>	<i>all other structure types</i>	Code, BibEntry, Em, FENote, Note, Quote, Reference, Span, Strong, Sub Lb1 as child of BibEntry, TOCI, or Note <i>pseudo items</i> Artifact, ASpan, ReversedChars, Clip
<i>part of the structure tree</i>	<i>yes</i>	<i>no</i>
<i>can extend across page boundaries</i>	<i>yes</i>	<i>no</i>
<i>can be interrupted by other elements</i>	<i>yes</i>	<i>no</i>
<i>can be activated with PDF_activate_item()</i>	<i>yes</i>	<i>no</i>
<i>nesting rules for child elements</i>	<i>may contain regular and direct elements</i>	<i>may contain only direct elements, and only if the parent can contain direct contain</i>

this rule in cases where an empty structure element conveys a semantic role in some other structure as in the following examples:

- ▶ Empty *TD* elements are required for empty table cells. This rule is honored by PDFlib's automatic table formatting (see Section 11.2.1, »Automatic Table Tagging«, page 278).
- ▶ Empty *LI* elements inside a list structure;
- ▶ Empty *Span* element to provide *ActualText* for white-space characters;
- ▶ Empty *Div* element to provide metadata properties or attributes;
- ▶ Empty *Document* element for a single-page document without any content;
- ▶ Empty *NonStruct* element to include arbitrary tagsets.

Programming scope and page boundaries. Most structure elements can only be created in *page* scope. Grouping elements may span multiple pages and can also be created in *document* scope. Pseudo items as well as direct items must be closed before ending or suspending a page. Tagging operations are not supported in template scope.

11.1.3 PDF 1.7 Structure Element Nesting Rules

In PDF 1.x a *Document* element should be used as root of the structure tree if the PDF contains a complete document. *Part or Sect* should be used as root element if the PDF contains a fragment of a document. The root element can conveniently be provided with the *tag* option of *PDF_begin_document()*. Table 11.3 lists the PDF 1.7 tag nesting rules for the structure element hierarchy.

Table 11.3 PDF 1.7 tag nesting rules for `PDF_begin_item()` and the tag/begintag options of various functions

item	rule
content items	<p>The following structure types may contain page content, i.e. text, image or vector graphics:</p> <p>H, H1, H2, ... P Lbl, LBody TH, TD Span, Quote, Note, Reference, BibEntry, Code Link, Annot Figure, Formula RB, RT, RP, WT, WP Artifact, ASpan, ReversedChars, Clip</p>
grouping elements	<p>Grouping elements must not contain content items, ASpan or ILSEs as children, i.e. a BLSE must be created before content can be created:</p> <p>Document, Part, Art, Sect, Div, BlockQuote, Caption, TOC, TOCI, Index, NonStruct, Private</p> <p>The option Placement=Block is recommended for the following elements as children of grouping elements: Figure, Formula, Form, Link, Annot.</p>
block-level elements	<p>The following block-level elements must not contain content items, i.e. a suitable grouping element or BLSE must be created before content can be created:</p> <p>L, LI, Table, TR, THead, TFoot, TBody</p> <p>Strict rule: the P element may not contain grouping elements.</p>
pseudo and inline elements	<p>Pseudo elements (i.e. Artifact, ASpan, ReversedChars, Clip) and the following ILSEs cannot have any descendants if direct=true:</p> <p>Code, BibEntry, Note, Quote, Reference, Span</p> <p>However, these elements may have children if direct=false.</p> <p>Artifacts cannot be created within pseudo elements and the following ILSEs with direct=true: Span, Quote, Note, Reference, BibEntry, Code</p>
table elements	<p>Table elements may contain one or more TR elements, or an optional THead followed by one or more TBody elements and an optional TFoot. TBody should not be the only child of Table. In addition, Table elements may have a Caption element as its first or last child.</p> <p>TR elements may contain TH and TD elements.</p> <p>TH and TD elements may not contain TR, TH, TD, THead, TBody, or TFoot</p> <p>THead, TBody, TFoot elements can contain only TR elements, and can only have Table as parent.</p> <p>TR may only have Table, THead, TBody, or TFoot as parent.</p>
list elements	<p>L elements may optionally contain a Caption element and one or more LI elements.</p> <p>LI elements may contain one or more Lbl or LBody elements or both. LI may only have L as parent. LBody may only have LI as parent.</p>
table of contents	<p>TOC elements may contain an optional Caption element as first child, and one or more TOCI and TOC elements (also in combination).</p> <p>TOCI elements may contain only Lbl, Reference, NonStruct, P, and TOC elements. TOCI may only have TOC as parent.</p>
label element	<p>Lbl may only have Annot, LI, Link, TOCI, BibEntry or Note as parent.</p>

Table 11.3 PDF 1.7 tag nesting rules for `PDF_begin_item()` and the tag/begin tag options of various functions

item	rule
interactive elements: links, form fields and annotations	<p>The following structure types are not allowed for the tag option of <code>PDF_create_field()</code> and <code>PDF_create_annotation()</code> or as parent for Link, Annot, and Form: table elements, ILSEs, Ruby and Warichu elements, pseudo elements.</p> <p>Annot elements cannot be nested.</p> <p>Link elements cannot be nested.</p> <p>Form elements must not contain any elements other than the OBJR element which is created by <code>PDF_create_field()</code> automatically.</p>
Japanese Ruby and Warichu	<p>Ruby can have RB, RT, and RP as children, but not any other structure types.</p> <p>Warichu can have WT and WP as children, but not any other structure types.</p>
PDF 2.0 consistency rules	<p>The following nesting rules for rarely used PDF 1.7 parent/child tag combinations have been introduced for consistency with PDF 2.0 nesting rules:</p> <p><i>Incompatible changes (stricter nesting rules):</i></p> <ul style="list-style-type: none"> ▶ RB, RT, and RP can only be children of Ruby ▶ WT and WP can only be children of Warichu ▶ Document may not contain Caption ▶ Caption may not contain Document or Caption ▶ H, H1 etc. may not contain Document, Part, or Div ▶ TH and TD may not contain Document or Caption ▶ Inline level elements with <code>direct=true</code> may not contain L, Table or Document

11.1.4 PDF 2.0 Structure Element Nesting Rules

PDF 2.0 tag nesting rules. PDF 2.0 imposes a strict regime of allowed vs. prohibited parent/child combinations. Table 11.4 lists PDF 2.0 tag nesting rules. The column »content items« specifies whether an element is allowed to contain page content (text, images and vector graphics as well as *OBJR* for interactive elements).

Table 11.4 PDF 2.0 tag nesting rules for `PDF_begin_item()` and the tag/begin tag options of various functions. Unless noted otherwise all children may appear an arbitrary number of times within the parent element.

parent element	allowed child elements	content items
structure tree root	The structure tree root must contain exactly one Document element. It may also contain a NonStruct element. The root can conveniently be provided with the tag option of <code>PDF_begin_document()</code> .	-
Document, DocumentFragment	Document, DocumentFragment, Part, Div, Sect, Aside, NonStruct, P, H1, ..., H ¹ , Title, Link, Annot, Form, L, Table, Figure, Formula, Artifact	-
Part	All structure types which are allowed as children of the Part element's first ancestor element which is different from Part.	-
Div	All structure types which are allowed as children of the Div element's first ancestor element which is different from Div.	-
Sect	DocumentFragment, Part, Div, Sect, Aside, NonStruct, P, H1, ..., H ¹ , Title, Lbl, Link, Annot, Form, FENote, L, Table, Caption, Figure, Formula, Artifact	-

Table 11.4 PDF 2.0 tag nesting rules for `PDF_begin_item()` and the tag/begin tag options of various functions. Unless noted otherwise all children may appear an arbitrary number of times within the parent element.

parent element	allowed child elements	content items
Aside	Document, DocumentFragment, Part, Div, Sect, NonStruct, P, H1, ..., H ¹ , Lbl, Link, Annot, Form, FENote, L, Table, Caption ¹ , Figure, Formula, Artifact	yes
NonStruct	All structure types which are allowed as children of the NonStruct element's first ancestor element which is different from NonStruct.	-
Title	Part, Div, Aside, NonStruct, P, Lbl, Em, Strong, Span, Link, Annot, Form, Ruby, Warichu, L, Table, Caption ¹ , Figure, Formula, Artifact	yes
	If the element is used as grouping element, the following element is also allowed as child, but it is mutually exclusive with content items: FENote	no
Sub	NonStruct, Lbl, Em, Strong, Span, Link, Annot, Form, Ruby, Warichu, L, Figure, Formula, Artifact	yes
P	NonStruct, Sub, Lbl, Em, Strong, Span, Link, Annot, Form, Ruby, Warichu, L, Figure, Formula, Artifact	yes
H1, ..., H	Sect ¹ , NonStruct, Sub, Lbl, Em, Strong, Span, Link, Annot, Form, Ruby, Warichu, Figure, Formula, Artifact	yes
Lbl	NonStruct, Sub, Em, Strong, Span, Link, Annot, Form, Ruby, Warichu, Figure, Formula, Artifact	yes
Em, Strong, Span	NonStruct, Sub, Lbl, Em, Strong, Span, Link, Annot, Form, Ruby, Warichu, Figure, Formula, Artifact	yes
Link	Div, Sect, NonStruct, Sub, Lbl, Em, Strong, Span, Annot, Ruby, Warichu, Figure, Formula, Artifact	yes
	If the element is used as grouping element, the following elements are also allowed as children, but they are mutually exclusive with content items: DocumentFragment, Part, Aside, P, Hn, H, Title, Form, FENote, L, Table, Caption	no
Annot	Div, Sect, NonStruct, Sub, Lbl, Em, Strong, Span, Link, Annot, Ruby, Warichu, Figure, Formula, Artifact	yes
	If the element is used as grouping element, the following elements are also allowed as children, but they are mutually exclusive with content items: DocumentFragment, Part, Aside, P, Hn, H, Title, Form, FENote, L, Table, Caption	no
Form	Div, NonStruct, Lbl, Caption ¹ , Artifact	yes
	If the element is used as grouping element, the following elements are also allowed as children, but they are mutually exclusive with content items: Part, L, Table, Figure, Formula	no
Ruby	NonStruct, RB, RT, RP	yes
Warichu	NonStruct, WT, WP	yes
RB, RT, RP, WT, WP	NonStruct, Sub, Em, Strong, Span, Link, Annot, Form, Artifact	yes
FENote	Part, Div, Sect, Aside, NonStruct, P, Sub, Lbl, Em, Strong, Span, Link, Annot, Form, Ruby, Warichu, L, Table, Figure, Formula, Artifact	yes
	If the element is used as grouping element, the following elements are also allowed as children, but they are mutually exclusive with content items: DocumentFragment, Caption	no

Table 11.4 PDF 2.0 tag nesting rules for `PDF_begin_item()` and the tag/begintag options of various functions. Unless noted otherwise all children may appear an arbitrary number of times within the parent element.

parent element	allowed child elements	content items
L	NonStruct, L, LI, Caption ¹ , Artifact	-
LI	Div, NonStruct, Lbl, FENote, LBody, Artifact	yes
LBody	Part, Div, Sect, Aside, NonStruct, P, H1, ..., H ¹ , Sub, Em, Strong, Span, Link, Annot, Form, Ruby, Warichu, FENote, L, Table, Caption ¹ , Figure, Formula, Artifact	yes
Table	NonStruct, TR, THead ¹ , TBody, TFoot ¹ , Caption ¹ , Artifact	-
TR	NonStruct, TH, TD, Artifact	-
TH, TD	Div, Sect, NonStruct, P, H1, ..., H ¹ , Lbl, Em, Strong, Span, Link, Annot, Form, Ruby, Warichu, FENote, L, Table, Figure, Formula, Artifact	yes
THead, TBody, TFoot	NonStruct, TR, Artifact	-
Caption	Part, Div, Sect, Aside, NonStruct, P, H1, ..., H ¹ , Sub, Lbl, Em, Strong, Span, Link, Annot, Form, Ruby, Warichu, FENote, L, Table, Figure, Formula, Artifact <i>If the element is used as grouping element, the following element is also allowed as child, but it is mutually exclusive with content items: DocumentFragment</i>	yes no
Figure	<i>If used as Block element:</i> Part, Div, Sect, Aside, NonStruct, P, H1, ..., H ¹ , Lbl, Em, Strong, Span, Link, Annot, Form, Ruby, Warichu, FENote, L, Table, Caption ¹ , Figure, Formula, Artifact	yes
Formula	Part, Div, Aside, NonStruct, P, H1, ..., H ¹ , Sub, Lbl, Em, Strong, Span, Link, Annot, Form, Ruby, Warichu, FENote, L, Table, Caption ¹ , Figure, Formula, Artifact <i>If the element is used as grouping element, the following element is also allowed as child, but it is mutually exclusive with content items: Sub</i>	yes no
Artifact	(All structure types; H and Caption may appear at most once.)	yes

1. This element may appear at most once within the parent element.

Additional PDF 2.0 tag nesting rules. In addition to the main tag nesting rules which restrict certain pairings of structure types, PDF 2.0 mandates additional restrictions beyond simple parent/child relationships according to Table 11.5. The mandatory PDF 2.0 requirements are enforced by PDFlib, while the recommended requirements are not enforced.

Table 11.5 Additional PDF 2.0 tag nesting rules for `PDF_begin_item()` and the tag/begintag options of various functions

structure type	requirement
The following mandatory PDF 2.0 rules are enforced by PDFlib:	
Document, DocumentFragment	Within each Document or DocumentFragment structure element, all heading structure elements must be Hn or H.
Annot	All annotations in an Annot structure element must be of the same annotation type. Annot must not be used for link annotations or form fields.
Form	Form fields must be included in a Form or Artifact element.

Table 11.5 Additional PDF 2.0 tag nesting rules for `PDF_begin_item()` and the `tag/begintag` options of various functions

structure type	requirement
Link	A Link element may contain only one link annotation, but not more.
Ruby	A Ruby element must contain one RB element followed by either an RT element or the sequence RP, RT, and RP.
Warichu	A Warichu element must contain the sequence WP, WT, and WP.
Caption	The Caption element must be the first or last element inside its parent element. The number of captions cannot exceed one.
H vs. H1, H2, ...	Either the unnumbered heading element H (strong structuring) or the numbered heading elements H1, H2, etc. (weak structuring) may be used with a Document or DocumentFragment element, but not a mixture of both.
The following recommended PDF 2.0 rules are not enforced by PDFlib:	
H	The H element should always be the first structure element within its parent element.
Title	The Title element should occur only once inside the parent grouping element, and it should occur at or near the beginning of the content inside that grouping element.
FENote	In order for text to be considered a footnote or endnote, there should be a reference from the enclosing content to the footnote or endnote. Such a reference may be achieved by means of a Link element through a structure destination in its link annotation or use of the <code>ref</code> option in a structure element.
Sub	If a Sub element is used, all other content inside the block level element that is the parent of the Sub element should also be enclosed in Sub structure elements.

11.1.5 Custom Structure Types, Role Maps and custom Tagsets

In addition to the standard structure types listed in Table 11.1 custom structure type names can be used. Custom structure types are typically used to localize structure type names (e.g. German *Abbildung* maps to *Figure*) or to work with application-specific type names (e.g. *Normal* maps to *P*). In order to facilitate repurposing documents which contain custom structure types the custom names must be mapped to their exact or approximate equivalent in the set of standard structure types. You can also remap standard structure types to other standard types in order to change their semantics. Custom elements types cannot be mapped to inline and pseudo elements. Element mapping for the standard types is achieved with the *rolemap* document option, e.g.

```
p.begin_document("tagged.pdf",
    "tagged=true lang=en rolemap={ {Heading H1} {Subhead H2} {Paragraph P} }");
```

Custom structure types must adhere to the nesting rules of the corresponding standard structure types to which they are role mapped.

Role map in Acrobat. You can view and edit the role map in Acrobat DC as follows:

- ▶ Select *View, Show/Hide, Navigation Panes, Tags*, click the menu button at the top of the *Tags* pane, and select *Edit Role Map* from the drop-down list.

Tagsets in PDF 2.0. PDF 2.0 supports two predefined tagsets which can be selected with `pdftagset=1.7` or `pdftagset=2.0`. In addition to these predefined tagsets custom tagsets can be defined. They are identified by a namespace in the form of a uniform re-

source identifier (URI). Custom tagsets must be declared in the *tagsets* document option. For each tagset the corresponding namespace URI and a role map must be provided, e.g.

```
tagsets={
  {namespace={http://www.pdflib.com/tagsets/foo} rolemap={{MyTag P} ...} }
  {namespace={http://www.pdflib.com/tagsets/bar} rolemap={...} }
}
```

Assuming a custom tagset has been declared this way the corresponding tags can be supplied to *PDF_begin_item()* and other tagging operations:

```
tag={tagname=MyTag namespace={http://www.pdflib.com/tagsets/foo}}
```

11.1.6 Artifacts

Relevant content and Artifacts. The contents of a page fall into one of the following categories:

- ▶ Relevant content has been created by the document author to convey the document's meaning. The document's logical structure tree describes the objects which comprise real content, and may also contain annotations.
- ▶ Graphic or text objects which do not contribute relevant page content, but have been created for pagination or layout purposes are called Artifacts. Artifacts are not included in the structure tree and are not read by a screen reader.

Marking Artifacts is strongly recommended to improve accessibility and is required in PDF/UA-1. Typical Artifacts are repeated headers and footers, page numbers, background images, and other items which are repeated on each page.

In PDF 1.7 Artifacts can only be expressed as direct elements, i.e. they are not included in the structure tree. In PDF 2.0 Artifacts can also be part of the structure tree (option *direct=false*). Unlike PDF 1.7 direct Artifacts, PDF 2.0 Artifacts in the structure tree have context (other entries in the structure tree). This is relevant e.g. for line numbers. Artifacts in the structure tree are also required for tagging annotations as Artifacts.

Artifacts in Acrobat. You can check Artifacts in Acrobat DC with one of the following methods:

- ▶ Select *Tools, Accessibility, Reading Order* to display or edit content elements on the page. Artifacts are called *Background* in Acrobat's *Reading OrderTool*. Unlike structure elements, they are not visualized with a frame and tag name when the tool is activated.
- ▶ Identify Artifacts by selecting *View, Show/Hide, Navigation Panes, Content*. The *Content* pane lists all page contents along with the respective structure type name or *Artifact* as appropriate. Because Artifacts are not read, there is no corresponding numbered block on the page. However, clicking on an Artifact in the list highlights the corresponding content element on the page.
- ▶ Search for Artifacts: select *View, Show/Hide, Navigation Panes, Tags*, click the menu button at the top of the *Tags* pane, *Find...*, and select *Artifacts* from the drop-down list. Since Artifacts are not part of the document structure, there is no corresponding entry in the *Tags* navigation pane nor the *Order* pane.
- ▶ Activate *View, Read Out Loud...* to let Acrobat read the structure elements on the page. Artifacts will not be read.

Designating content as Artifact. Artifacts can be specified in PDFlib with the *Artifact* tag name in `PDF_begin_item()`:

```
id = p.begin_item("Artifact", "");
```

Alternatively, Artifacts can be specified with abbreviated tagging, i.e. with the *tag* option of various functions:

```
p.fit_textline(text, x, y, "tag={tagname=Artifact}");
```

It is recommended to create Artifacts only if no BLSE is currently active. However, since this may not always be possible for applications, PDFlib automatically interrupts the currently active element when an Artifact is created and activates it again after the Artifact. Note that the tag nesting rules (see »PDF 1.7 Structure Element Nesting Rules«, page 266) don't allow Artifacts within ILSEs.

Classifying Artifacts. Irrelevant page content should be identified with the *Artifact* pseudo tag, and classified according to one of the following keywords of the *artifacctype* option:

- ▶ *Pagination* Artifacts: page features such as running heads and page numbers. Pagination Artifacts can be further classified with the *artifactssubtype* option and one of the keywords *Header*, *Footer*, *Watermark*.
- ▶ *Layout* Artifacts: typographic or design elements such as rules and table shadings;
- ▶ *Page* Artifacts: production aids, such as trim marks and color bars;
- ▶ *Background* Artifacts: images or colored areas that run the entire width or height of the page or the entire dimensions of a structure element.

Pagination and Background Artifacts support the *Attached* option which specifies to which page edge or edges the Artifact is attached (*Top/Bottom/Left/Right*).

The following example creates a pagination Artifact with subtype *Header*:

```
id = p.begin_item("Artifact",  
    "artifacctype=Pagination artifactssubtype=Header Attached={Top Left}");
```

Automatic Artifact tagging. Since all page contents should be tagged either as structure element or as Artifact PDFlib automatically tags certain graphical elements. The following decorative elements are automatically tagged as *Artifact* with *artifacctype=Layout* in Tagged PDF mode:

- ▶ Block decoration: all decorative elements created by `PDF_fill_block()`, i.e. stroking and filling created according to the *backgroundcolor* and *bordercolor* properties.
- ▶ Matchbox decoration: matchbox rectangles created according to the matchbox options *fillcolor*, *shading*, and *strokecolor*
- ▶ Table decoration: if automatic table tagging is active (see Section 11.2.1, »Automatic Table Tagging«, page 278), table ruling and shading, i.e. stroking and filling according to the options *fill*, *stroke*, *showborder*, *showgrid*. In the second and subsequent table instances all rows created due to the options *header*, *footer*, and *caption* are also marked as Artifact.
- ▶ Textline Artifacts: *leader*, *shadow*, *showborder*
- ▶ Textflow Artifacts: *leader*, *shadow*, *showborder*, *showtabs*
- ▶ Text decoration: *underline*, *overline*, *strikeout*

Like all Tagged PDF features, automatic Artifact tagging works only in *page* scope.

11.1.7 Text Handling

Language specification. In Tagged PDF the natural language of text should be specified explicitly; this allows screen readers to switch to the appropriate language when reading the document. The natural language can be specified on different levels:

- ▶ The *lang* option in *PDF_begin_document()* should be set to specify the primary language, i.e. the natural language of the document as a whole. This specification covers page contents as well as interactive elements such as bookmarks and annotations. The language specification may be changed in *PDF_end_document()*. This may be useful e.g. if the language is chosen based on an imported PDF document.
- ▶ The document language can be overridden for individual structure elements with the *lang* option in *PDF_begin_item()* or the *tag* option of various functions.
- ▶ In PDF 2.0 the *lang* option can also be supplied to *PDF_create_annotation()* to specify the natural language for annotations.

Contents which are encoded as text but are not part of a natural language, such as programming code, musical notes, type samples or mathematical notation, should use an empty language code, e.g. *lang={}*.

Hypertext strings may include Unicode language identifiers for specifying the language. However, these are not supported in current PDF viewers. Unicode language identifiers are comprised of the following sequence:

- ▶ The Unicode Escape character U+001B (two bytes).
- ▶ An ISO 639 language code (two ASCII bytes), e.g. *en, ja, de*.
- ▶ Optionally an ISO 3166 country code (two ASCII bytes), e.g. *US, JP*.
- ▶ The Unicode Escape character U+001B (two bytes).

Examples in hexadecimal notation:

```
001B656E5553001B      (=enUS)
001B7A68001B          (=zh)
001B6465001B          (=de)
```

In Unicode-capable languages the two ASCII characters must be packed to form a single Unicode value:

```
\u001B\u6465\u001B      (=de)
```

In the C language the two ASCII characters must be packed into a single Unicode character. With *charref=true* the sequence can be expressed as follows:

```
&#x001B;&#x6465;&#x001B;    (=de)
```

Separate words with space characters. Words should be separated by space characters (U+0020). The *autospace* option can be used for automatically generating space characters after each call to one of the text output functions.

Hyphenation. Hyphenation, i.e. splitting a word in two parts at the end of a line, must be represented using a soft hyphen character (U+00AD) as opposed to a hard hyphen or minus character (U+002D). The soft hyphen character ensures that Acrobat can correctly recombine (dehyphenate) the hyphenated word when searching for text. When text contains a soft hyphen U+00AD the text engine uses the glyph for U+00AD if it is available in the font, and U+002D otherwise. If the font contains two separate glyphs for

U+00AD and U+002D using the soft hyphen U+00AD in the text is sufficient for fulfilling Tagged PDF requirements for hyphenation. Use the following code fragment to check whether a font supports this simple hyphenation method:

```
gid1 = (int) p.info_font(font, "glyphid", "unicode=0x00AD"); /* soft hyphen */
gid2 = (int) p.info_font(font, "glyphid", "unicode=0x002D"); /* hyphen-minus */

if (gid1 != -1 && gid2 != -1 && gid1 != gid2)
    System.err.println("font contains separate soft hyphen glyph");
else
    System.err.println("font doesn't contain separate soft hyphen glyph");
```

If the font doesn't contain a separate glyph for U+00AD (or another hyphenation character is used), the hyphen must be tagged as *Span* or *ASpan* with an *ActualText* attribute containing U+00AD (note that this is not possible in PDF 1.4). When multiline text is created with Textflow the *ActualText* is assigned automatically to the specified *hyphenchar* (and *autospace* is suppressed).

Attaching the required *ActualText* for Textlines can be achieved as follows:

- ▶ In *PDF_fit_textline()* the option *tagtrailinghyphen* results in suitable *ActualText* and suppressed *autospace*. Since the option defaults to U+00AD proper tagging happens automatically if the text uses U+00AD as hyphen.
- ▶ If a font does not contain a separate glyph for U+00AD you can amend it with the soft hyphen from a suitable fallback font using the following font loading option:

```
fallbackfonts={{fontname=AuxiliaryFont forcechars=x00AD}}
```

- ▶ Suitable *ActualText* can be assigned manually as follows; keep in mind that this is only required if the font doesn't contain two separate glyphs for U+00AD and U+002D (this code works only with PDF 1.5 or above; add the option *direct=false* for PDF 1.4):

```
p.set_option("charref=true");
p.fit_textline("-", x, y, "tag={tagname=ASpan ActualText=&#x00AD;}");
```

11.1.8 Alternate Description, Replacement Text and Abbreviation Expansion

Tagged PDF provides features which enhance accessibility of images and text which cannot easily be read without additional information.

Alternate Description (Alt). Items which do not translate naturally to text can be assigned an alternate description via the *Alt* option. Examples are images, formulas, and annotations without the *contents* option.

The alternate description should consist of a whole word or phrase which can be read by a screen reader. The description should end with a period or space character as appropriate to ensure that screen readers don't merge it with subsequent text. It is recommended to avoid initial phrases such as »*This image shows...*« in the alternate description. The *Alt* value provides a description of the structure element and all of its children. The *ASpan* pseudo element can be used to assign *Alt* to some part of a structure element.

For example, the image of a company logo could be described via the *Alt* option:

```
p.fit_image(image, x, y, "tag={tagname=Figure Alt={Kraxi company logo }}")
```

Replacement text (ActualText). Items which resemble text, but where the text is represented in some non-standard way can be assigned equivalent replacement text via the *ActualText* tagging option. Examples are illustrations showing swash characters or drop caps, and images which use pixels to represent a word. On the other hand, OCR text for a scanned page should not be supplied as *ActualText*, but rather as invisible text (i.e. *textrendering=3*).

The replacement text should contain one or more characters which are equivalent to what a person would see when viewing the content. The *ActualText* value serves as replacement for the structure element and all of its children. The *ASpan* pseudo element can be used to assign *ActualText* to some part of a structure element.

For example, the symbolic *flower* glyph for which no Unicode value is available can be assigned suitable *ActualText* to make it clear that this glyph is actually used as a bullet character U+2022:

```
p.fit_textline("&.flower;", x, y, "tag={tagname=ASpan ActualText={&#x2022;} } ...");
```

Nesting rules for alternate and replacement text. The following rules must be obeyed when using the *Alt* and *ActualText* attributes:

- ▶ Since *Alt* and *ActualText* cover the full sub-hierarchy below the affected structure element, both options are not allowed for a structure element if any of its ancestors in the structure hierarchy already contains an *ActualText* or *Alt* attribute.
- ▶ An element with an *Alt* or *ActualText* attribute must include content items or one or more *Link* elements. If the attribute is applied to a *Link* element it must contain content items or one or more *OBJR* references created by *PDF.create_annotation()* (see »Links and other annotation types«, page 282). Otherwise it would be impossible to determine the page on which the attribute should be read. This rule applies to the element itself; it is not sufficient to have child elements with content items.

Alternate and replacement text in Acrobat. You can display the *Alt* and *ActualText* attributes of a structure element with Acrobat DC as follows:

- ▶ Select *View, Show/Hide, Navigation Panes, Tags*, right-click on a structure element in the hierarchy and select *Properties...* to display the object properties dialog. The *Tag* tab displays the *Actual Text* and *Alternate Text* attributes.

Abbreviation expansion (E). Abbreviations and acronyms can be assigned expansion text with the *E* suboption of the *tag* option. The expansion text should consist of a whole word or phrase which can be read by a screen reader. If an abbreviation doesn't have any expansion text the *E* attribute can be supplied nevertheless in order to assist the text-to-speech conversion process. For example, the term *IBM* may have expansion text *I B M* (with intermediate space characters) assigned to it.

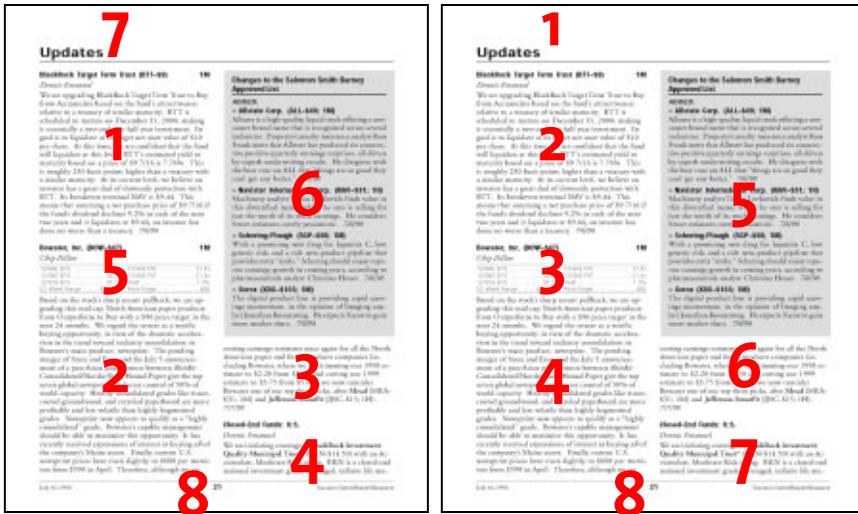
In the following code fragment the expansion text *Mister* is assigned to the abbreviation *Mr.*:

```
p.fit_textline("Mr.", x, y, "tag={tagname=ASpan E={Mister} } ...");
```

11.1.9 Print Stream Order and Logical Reading Order

There are two fundamentally different concepts for content ordering in PDF. figure 11.2 visualizes a sample page with two main text columns, interrupted by a table and an inserted summary on a gray background as well as header and footer.

Fig. 11.2
Logical reading order
(left) and print stream
order (right).



The order of PDFlib API method calls determines the order of PDF text and drawing operators in the page content stream (called »raw print stream order« in Acrobat). Since page contents may be created in any way which suits the controlling application this is only a technical ordering which doesn't necessarily have any semantic significance. The print stream order is displayed in Acrobat's *Order* and *Contents* panes.

Logical reading order is the order in which a human reads the text. It determines the ordering used by a screen reader and Acrobat's *Read Aloud* feature. The logical reading order is determined by the logical structure tree. Tagged PDF requires that all content is tagged in semantically correct order, i.e. the structure hierarchy must include the page contents in the order in which they will be read by a human. Proper tagging ensures that screen readers present the contents in logical order. Since Artifacts are not part of the structure tree they are excluded from the logical reading order.

Reading order and print stream order in Acrobat. You can check the logical reading order in Acrobat DC with the following methods:

- ▶ Select *View, Show/Hide, Navigation Panes, Tags* and check the order of elements from top to bottom. This ordering should accurately reflect the desired reading order.
- ▶ Activate *View, Read Out Loud...* to have Acrobat read the page contents in the reading order specified in the document.

The *Order* and *Content* panes list page contents in print stream order.

Creating content in logical reading order. The natural method which works in many situations is to sequentially generate all constituent parts of a structure element, and then move on to the next element in the logical sequence. In technical terms, the structure tree is created by a single depth-first traversal, and PDFlib functions to create page contents are called in the order in which the contents will be read.

PDFlib supports several methods for creating contents in any order while still creating the structure hierarchy in logical order. These methods are discussed in Section 11.2.4, »Creating Contents out of Order«, page 288.

11.1.1 Tagged PDF Problems in Adobe Acrobat

This section mentions observations that we made while testing Tagged PDF output in Adobe Acrobat. Table 11.6 lists bugs and inconsistent behavior in Acrobat, grouped according to the following features:

- ▶ Acrobat's accessibility checker: Acrobat's accessibility checker can be used to determine the suitability of Tagged PDF documents for consumption with assisting technology such as a screen reader.
- ▶ Acrobat's Read Aloud Feature: Tagged PDF enhances Acrobat's capability to read text aloud.
- ▶ The »Find...« function in Acrobat's Tags pane can be used to search for Artifacts and unmarked content.

Table 11.6 Acrobat DC problems related to Tagged PDF

Description, recommendations and workarounds
Acrobat's accessibility checker
The Alt attribute is ignored for Figure tags.
The accessibility checker complains about form fields of type check box with a check mark symbol »Character encoding – failed« although Read aloud perfectly reads the field contents.
Acrobat's Read Aloud feature
When a tagged page has been placed with PDI and contains only Artifacts, Read Aloud nevertheless reads the contents of the placed page.
Other Acrobat functions
The Find... function in the Tags pane incorrectly reports Artifacts as unmarked content, e.g. table decoration. In contrast, these items are correctly displayed as »Container <Artifact>« in the Content pane.

11.2 Advanced Tagged PDF Creation Topics

11.2.1 Automatic Table Tagging

`PDF_fit_table()` can automatically create suitable tags for the generated table based on the information supplied to the `PDF_add_table_cell()` calls for the table contents. The `tag` option of `PDF_fit_table()` with `tagname=Table` triggers automatic table tagging as detailed in Table 11.7.

Note Automatic table tagging can only be leveraged if PDFlib's table engine is used. While it is possible to correctly tag tables which are created manually (i.e. without PDFlib's table engine), this process requires detailed knowledge of the table structure in the client application. In addition to the row/column structure relevant information about header cells and row/column spans is required. Empty cells must also be tagged.

Visualizing table tags in Acrobat. You can visualize the structure of table elements with Acrobat DC as follows:

- ▶ Select *Tools, Accessibility, Reading Order*. Table elements are highlighted and marked with a small number near the top left corner of the table. If a table summary is present, it is displayed next to the small number.

Table 11.7 The tag option in `PDF_fit_table()` and automatic table tagging

tag option in <code>PDF_fit_table()</code>	result
<code>tagname=Table</code>	Recommended approach: activates automatic table tagging. Instead of <code>tagname=Table</code> a custom tag which is role-mapped to <code>Table</code> can also be used.
<code>tagname=Artifact</code>	The whole table contents including the caption are marked as <code>Artifact</code> ; the <code>BBox</code> attribute is added automatically.
any other <code>tagname</code>	No <code>Table</code> structure is created, but the cell contents are added as children of the element specified in the tag option. Pseudo and inline items are not allowed for <code>PDF_fit_table()</code> .
tag option not supplied	No automatic table tagging. If the tag option is supplied to individual calls to <code>PDF_add_table_cell()</code> (as suboption to one of the <code>fit*</code> options) the corresponding structure elements for the cell contents are created, but without any table structure. This may be useful if the table is used for layout purposes instead of as a data table.

Travel Expense Report				
	Meals	Hotels	Transport	subtotals
Madrid				
3-Aug-2012	37.74	112.00	45.00	
4-Aug-2012	27.28	112.00	45.00	
subtotals	65.02	224.02	90.00	379.02
Paris				
5-Aug-2012	96.25	109.00	36.00	
6-Aug-2012	35.00	109.00	36.00	
subtotals	131.25	218.00	72.00	421.25
Totals	196.27	442.00	162.00	827.27

Fig. 11.3

Acrobat's Table Editor for Tagged PDF tables displays header (TH) and data (TD) cells.

- ▶ Select the number or structure type name in the top left corner of the table and click *Table Editor* in the *Reading Order* dialog. The table structure is visualized with horizontal and vertical lines. In the *Table Editor Options* dialog you can instruct Acrobat to display *TH/TD* icons according to the type of table cell (see figure 11.3).
- ▶ Right-click on a table cell and select *Table Cell Properties...* to check the cell type (Header cell *TH* vs. data cell *TD*), the *scope* attribute, *rowspan* and *colspan* attributes and *header/ID* values.

Note the following restrictions when working with the Table Editor in Acrobat:

- ▶ The lines which visualize table cells may be displayed in wrong positions.
- ▶ Acrobat does not activate the Table Editor if the table uses a custom structure element which is role mapped to *Table* (as opposed to the standard element *Table*).
- ▶ The Table Editor does not display *Caption* elements which may be present in a table.
- ▶ If a table cell contains vertical text (e.g. a *Textline* with *orientate=east* or *west*) this cell and its immediate neighbor to the right are not displayed in the Table Editor although they are present in the logical structure tree and their contents are visible on the page.

Automatically created table tags and attributes. Automatic table tagging works only in page scope and with *tag*={*tagname*=Table} in *PDF_fit_table()*. The detailed operation is as follows:

- ▶ A *Table* element is created for each logical table. If a table is split in two or more instances only a single *Table* element is created. The *Summary* attribute is added to the *Table* element if the *Summary* suboption has been supplied to the *tag* option of *PDF_fit_table()*.
- ▶ A *Caption* element is created in the first table instance if the *caption* option was specified in *PDF_fit_table()*. As a grouping element *Caption* does not allow any content items as children. You must therefore supply the *tag* suboption of the *caption* option to specify a structure element as child of *Caption*. This element can hold the actual contents of the caption. Captions created in the second and subsequent table instances are marked as *Artifact* since only a single *Caption* element is allowed for each *Table* element.
- ▶ A *TR* element is created for each table row. In the first table instance rows which are specified in the *header* option of *PDF_fit_table()* are wrapped by *TH*Head, rows which are specified in the *footer* option are wrapped by *TF*Foot. All other rows are wrapped by *T*Body if headers or footers are present. In the second and subsequent table instances header and footer rows are marked as *Artifact* since the header and footer row groups must only appear once.
- ▶ Each table cell is wrapped in a *TH* (table header) or *TD* (table data) element according to the *tagname* suboption of the *tag* option of *PDF_add_table_cell()*. If this option has not been supplied the cell type is selected as follows:
 - ▶ The *Scope* attribute for a cell enforces *TH* (even if *tagname*=*TD* is specified).
 - ▶ If another cell contains the *tag* option with the *Headers* suboption with the *id* of the current cell, the current cell is forced to *TH* (even if *tagname*=*TD* is specified).
 - ▶ If the cell is included in a table row which is part of the table header as specified by the *header* option of *PDF_fit_table()* it is wrapped by *TH*, and *Scope*=*Column* is added.
- ▶ An empty *TD* element is created for each table cell for which *PDF_add_table_cell()* has not been called.
- ▶ *TH* and *TD* elements get appropriate *RowSpan* and *ColSpan* attributes according to the *rowspan* and *colspan* options of *PDF_add_table_cell()*. The *RowSpan* and *ColSpan* suboptions of the *tag* option cannot be used.
- ▶ *Table*, *TH* and *TD* elements are assigned appropriate *Width* and *Height* attributes; *Table* elements are also assigned the *BBox* attribute.
- ▶ Other table cell attributes may be supplied as suboptions for the *tag* option of *PDF_add_table_cell()*. The following options are not allowed: *RowSpan*, *ColSpan*, *Height*, *index*, *parent*, *Width*.
- ▶ Table rows and cells are emitted in zigzag order starting from the top left cell (i.e. column 1, row 1) to the bottom right cell, regardless of the order of calls to *PDF_add_table_cell()*.
- ▶ Decorative table elements are automatically tagged as *Artifact* with *artifacttype*=*Layout*, i.e. ruling and shading (fill/stroke) of table cells, rows, columns or the full table, matchbox filling and ruling, *showborder* rules, and visualization aids controlled via *debugshow*, *showcells*, and *showgrid*.
- ▶ The *tag* option is ignored in the second and subsequent calls to *PDF_fit_table()* for the same table since automatic tagging takes care of all relevant tags.

Cookbook Code samples for automatic table tagging can be found in the `table_pdfua1` and `invoice_pdfua1` topics in the `pdfua` category of the *PDFlib Cookbook*.

Adding tags and attributes to table cells. Abbreviated tagging can be applied to the table caption, a table cell, or the contents of a table cell. Supplying the `tag` option to `PDF_add_table_cell()` is useful in the following situations:

- ▶ Table cells can be forced to be header cells (instead of *TD* data cells) with `tagname=TH` if they are not included in header rows or don't have any `Scope` attribute.
- ▶ Create the proper tag structure required for links; see »Tagging links in a table cell«, page 283, for details.

Some restrictions apply to the suboptions of the `tag` option in `PDF_add_table_cell()`:

- ▶ The following options cannot be used: `ColSpan`, `Height`, `index`, `parent`, `RowSpan`, `Width`.
- ▶ The `tagname` option can only have the value `TH` or `TD` to specify the type of table cell. However, descendant tags can be specified by nesting the `tag` option.

Adding tags and attributes to table cell contents. You can also supply `tag` as suboption to the following options (or corresponding suboptions of the `caption` option) of `PDF_add_table_cell()`:

`fitannotation`, `fitfield`, `fitgraphics`, `fitimage`, `fitpath`,
`fitpdipage`, `fittextflow`, `fittextline`

This is useful in the following situations:

- ▶ Specify sub-structure for the contents of a table cell. The `tag` option creates a child element of the cell's *TH* or *TD* element. The following values for `tagname` are not allowed for the `tag` option if automatic table tagging is active (in other words, nested tables are not supported):

`Table`, `TR`, `TH`, `TD`, `THead`, `TBody`, `TFoot`

- ▶ The table attributes listed below cannot be created automatically, but may be required for accessibility purposes. They must be supplied by the user:
 - ▶ The `Headers` option must be supplied for *TD* cells which reference one or more *TH* cells inside the table (i.e. they reference header cells which are not part of the `header` option of `PDF_fit_table()`).
 - ▶ The `Id` and `Scope=Row` options must be supplied for *TH* cells which are referenced in any `Headers` option (`Scope=Column` is created automatically for *TH* column header cells).
- ▶ The caption may contain arbitrary contents which can itself be tagged. For example, the following option list creates a `Caption` element containing a single `Textline` within a nested `P` element:

```
caption={ fittextline={tag={tagname=P title={Travel Expense Report}} ... } ... }
```

11.2.1 Tagging Interactive Elements

Links, annotations and form fields must also be made accessible. The corresponding interactive elements must be represented in the structure tree at the proper location.

Cookbook A code sample for creating tagged links can be found in the `starter_pdfua1` sample. The `image_with_link_pdfua1` topic creates a link with a background image. The

table_of_contents_pdfua1 *topic creates a table of contents with TOC and TOCI structure elements and actionable links.*

Links and other annotation types. Annotations require the following items for accessibility (see figure 11.4):

- ▶ A *Link* element (for link annotations) or *Annot* element (for all other annotation types) serves as container for the next two items. *Alt* or *ActualText* options may be provided to supply an alternate description or replacement text. The *Alt* attribute of a *Link* element should describe the purpose of the link. If the link target is located in the current document (*GoTo* action) the *Link* element should additionally be contained in a *Reference* element (with option *direct=false*).
- ▶ Text which represents the interactive element should be created inside the container element. If no text is required this element can be skipped, e.g. for a Text annotation. If the annotation is represented by a raster image or vector graphics this should be tagged as *Artifact*. In this case the alternative text of the link describes both the graphic and the link. It is recommended to use the *matchbox* option of various content creation functions to prepare geometry information for the next element. The currently active element when the *matchbox* option is supplied must be a structure item, i.e. the option *direct=true* is not allowed.
- ▶ The annotation must be created with *PDF_create_annotation()*. In addition to the annotation this method creates a corresponding structure element of type *OBJR* (object reference) for the annotation. The *contents* option of *PDF_create_annotation()* should be provided for link annotations (this is required in PDF/UA-1). Other annotation types should be created with the *contents* option of *PDF_create_annotation()* or the *ActualText* tagging option. The *usematchbox* option can be used to conveniently supply the geometry of the visual content created in the second step.

The second and third items may be created in either order. The annotation requirements above also apply to annotations created in table cells with the option *fitannotation* of *PDF_add_table_cell()*.

Explicitly tagging links. The following code fragment creates an interactive link with the required three items (the resulting tag structure is shown in figure 11.4):

```
/* Create the parent Link element */
id_link = p.begin_item("Link", "Title={Kraxi on the Web} Alt={Kraxi on the Web}");

/* Create visible content which represents the link */
p.fit_textline("Click here to go to the Kraxi website", x, y,
    "matchbox={name={kraxi}} fontsize=14 font=" + font);

/* Create URI action */
action = p.create_action("URI", "url={http://www.kraxi.com}");

/* Create Link annotation on named matchbox "kraxi". */
p.create_annotation(0, 0, 0, 0, "Link",
    "action={activate=" + action + " } "
    "usematchbox={kraxi} contents={Link to Kraxi Inc. Web site}");

p.end_item(id_link);
```

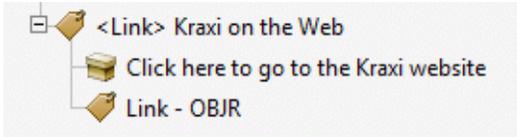


Fig. 11.4
Representation of a link in the structure tree

Tagging links with tags in Textflow option lists. Correct tagging of links inside a Textflow can be achieved with the matchbox feature and tags in Textflow option lists. Using the general idiom shown in »Adding a Web link to the Textflow matchbox«, page 242. you can add suitable tags for Link annotations with the code fragment below. Textflow automatically creates as many matchbox rectangles as required, so a single call to `PDF_create_annotation()` is sufficient for creating and tagging all instances of a particular named matchbox on a page:

```
final String text =
"Visit our <tag={tagname=Link} matchbox={name=kraxi}>Web site<tagend matchbox=end>";
...
tflow = p.create_textflow(text, optlist);
result = p.fit_textflow(tflow, 0, 0, 50, 70, "fitmethod=auto");
...
/* create URI action */
act = p.create_action("URI", "url={http://www.kraxi.com}");

/* create Link annotation on named matchbox(es) "kraxi" */
p.create_annotation(0, 0, 0, 0, "Link",
    "action={activate " + act + "} " +
    "usematchbox={kraxi} contents={Link to Kraxi Inc. Web site}");
```

A matchbox may be comprised of multiple rectangles, e.g. due to hyphenation or because multiple instances of the *matchbox* option supplied the same name. In this situation PDFlib by default creates a single Link annotation consisting of multiple active areas (one for each rectangle of the matchbox). Alternatively you can use the option `singlelink=false` of `PDF_create_annotation()` to create multiple link annotations where each annotation consists of a single rectangle. Note that this is not allowed in PDF 2.0.

Cookbook A code sample for creating Tagged links in a Textflow can be found in the topic `textflow_pdfua1` in the `pdfua` category of the PDFlib Cookbook.

Tagging links in a table cell. Links in a table cell require the tag structure described above. However, this can be a bit tricky because the structure elements *TH/TD*, *Link*, *OBJR* and possibly content must be properly nested. In order to achieve this you must supply the *tag* option to `PDF_add_table_cell()` and make use of its nesting feature. The option list below fills a table cell with a line of text and a link annotation. The enclosing *TD* element is supplied in the outer *tag* option (since *TD* is created automatically by the table engine the outer *tagname* suboption can be omitted), and the *Link* element is supplied in the inner *tag* option. Finally, the required *OBJR* element is created automatically by the *fitannotation* option which serves as equivalent to `PDF_create_annotation()`:

```
fittextline={font=... fontsize=25 fillcolor=blue}
annotationtype=Link fitannotation={contents={Kraxi home page} action={activate ...}}
tag={tagname=TD tag={tagname=Link}}
```

Tagging form fields. Form fields require the following items for accessibility (see figure 11.5):

- ▶ A group of related form fields may be grouped with a *Part* structure element.
- ▶ A *Div* structure element is recommended as container of the following elements of an individual form field.
- ▶ A *Caption* structure element with a *P* structure element that encloses text which describes the purpose of the field.
- ▶ For checkboxes and radio buttons it is recommended to create another *Div* element which encloses the check boxes/radio buttons and their corresponding captions. Since the *Caption* structure element logically usually precedes the field it comes before the *Form* element in the structure hierarchy, even if they are placed in reverse order on the page.
- ▶ The form field must be created with `PDF_create_field()`. It should be contained in a *P* element which contains a *Form* element. In addition to the field this method creates a structure element of type *OBJR* (object reference) inside the *Form* element. *Alt* or *ActualText* options may be provided to supply an alternate description or replacement text. The *tooltip* option of `PDF_create_field()` should be provided to enhance accessibility of the field (this is required in PDF/UA-1). When creating radio buttons, no *Form* element is required in `PDF_create_fieldgroup()` for the radio button group, but only in `PDF_create_field()` for creating the individual buttons.

The following code fragment creates the recommended structure for a text field as shown in figure 11.5. The *Caption* and *P* structure elements are created with nested abbreviated tagging using the *tag* option in `PDF_fit_textline()`. Abbreviated tagging in `PDF_create_field()` is used to create the nested *P* and *Form* structure elements:

```
id_Div = p.begin_item("Div", "");

labeltext = "Enter name:";
p.fit_textline(labeltext, x1, y1, "tag={tagname=Caption tag={tagname=P}} ...");

p.create_field(x2, y2, x3, y3, "name", "textfield",
    "tag={tagname=P tag={tagname=Form}} tooltip=\"" + labeltext + "\" +
    "bordercolor={gray 0} font=" + font);

p.end_item(id_Div);
```

The form field requirements above also apply to form fields created in table cells with the option `fitfield` of `PDF_add_table_cell()`.

Cookbook A code sample for creating all types of Tagged form fields can be found in the topic `form_fields_pdfua1` in the `pdfua` category of the *PDFlib Cookbook*.

Structured bookmarks. Bookmarks can be assigned a structure element in addition to the usual page-oriented destination. Such bookmarks are called *structured bookmarks* and Acrobat offers additional features for them. Upon right-clicking a structured bookmark in Acrobat the functions *Delete Page(s)* and *Extract Page(s)* are available which operate on the page or pages which contain the structure element. Structured bookmarks create a connection between a bookmark and a structure element. This connection can be created in two different ways:

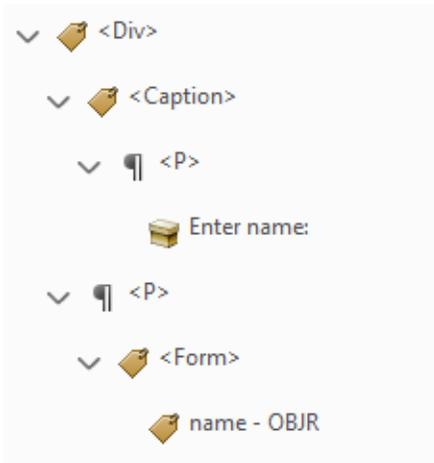


Fig. 11.5
Representation of a form field in the structure tree

- ▶ Create a bookmark with `PDF_create_bookmark()` and supply its handle to the `bookmark` option in `PDF_begin_item()` or the `tag` option of various functions:

```
bm = p.create_bookmark("Section 1", "");
id = p.begin_item("H1", "Title={Section 1} bookmark=" + bm);
p.fit_textline(text, x, y, "");
p.end_item(id);
```

This method can also be used with abbreviated tagging:

```
bm = p.create_bookmark("Section 1", "");
p.fit_textline(text, x, y,
"tag={tagname=H1 Title={Section 1} bookmark=" + bm + ""});
```

The disadvantage of this method is that the bookmark text must be available before the structure element and its contents are created. This may be inconvenient if the referenced structure element is located higher up in the structure tree.

- ▶ Create a structure item with `PDF_begin_item()` and supply its handle to the `item` option of `PDF_create_bookmark()`. Instead of an item handle the keyword `current` can be supplied as a shortcut which refers to the current structure element at the time `PDF_create_bookmark()` is called:

```
id = p.begin_item("H1", "Title={Section 1} ");
bm = p.create_bookmark("Section 1", "item=current");
p.fit_textline(text, x, y, "");
p.end_item(id);
```

This method has the advantage that the bookmark text must only be available when the structure element and its contents are created. However, it cannot be used with abbreviated tagging.

Structured bookmarks can only refer to open structure elements, but cannot refer to pseudo or inline items. The client code must ensure that the bookmark's destination matches the structure element (otherwise activating the bookmark in Acrobat would not jump to the element, but to a different location in the document). If the associated

structure element spans more than one page the bookmark should point to the first page in this range.

11.2.2 PDF 2.0 Features for Tagging Interactive Elements

Interactive elements as Artifacts. Interactive elements cannot be classified as Artifacts in PDF 1.7 because Artifacts can only be expressed as marked content items. In PDF 2.0 Artifacts can be included in the structure tree, which can be used to tag interactive elements as Artifacts. The following code fragment (which requires *compatibility=2.0*) creates a Link annotation as Artifact:

```
p.create_annotation(100, 100, 300, 200, "Link",  
    "tag={tagname=Artifact direct=false}");
```

Structure destinations. In PDF 1.7 destinations point to a geometric location on a particular page in the document. Structure destinations, a PDF 2.0 feature, point to some structure element in the document hierarchy and are therefore independent of physical locations. This technique requires two steps:

- ▶ The *elementid* tagging suboption assigns a unique string to the structure element which serves as destination, e.g. for a Link or bookmark:

```
p.fit_textline("Introduction",  
    50, 700, "tag={tagname=H1 elementid={Intro}} fontsize=24");
```

- ▶ A *GoTo* action uses the designated structure element as destination:

```
action = p.create_action("GoTo", "destination={type=fitwindow structdest={Intro}}");
```

The element id may be assigned to the target element before or after use. Structure destinations are not limited to *GoTo* actions, but can be used in all situations where regular destinations are possible, e.g. bookmarks and annotations. Structure destinations can also be used to define named destinations.

11.2.3 Tagging Lists

Lists are used to group related items. They are represented by the following structure elements:

- ▶ An *L* element which contains all of the following structure elements. The *List-Numbering* option can be used to specify the numbering system used in the *Lbl* elements. The *ListNumbering* option may be useful for screen readers even without *Lbl* elements.
- ▶ An optional *Caption* element. Since *Caption* is a grouping element it cannot contain content items, but only other structure elements (e.g. *P*).
- ▶ One or more list items (*LI*) containing the following:
 - ▶ An optional label (*Lbl*) with a bullet, number, etc.
 - ▶ An *LBody* element with the actual contents of the list item. *LBody* may contain either content items or other structure elements including a nested list.

Explicitly tagging lists. The following code fragment creates a list with a caption and three items. Each list item is preceded by a bullet character U+2022 which is tagged as label (the resulting tag structure is shown in figure 11.6):

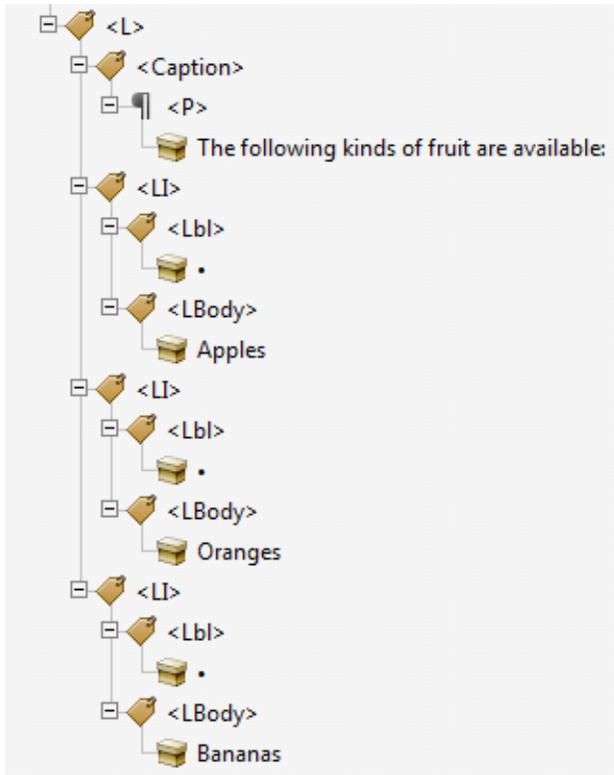


Fig. 11.6
Representation of an accessible list
in the structure tree

```
id_list = p.begin_item("L", "ListNumbering=Disc");

/* Create both Caption and P elements at once */
p.fit_textline("The following kinds of fruit are available:",
    x1, y, "tag={tagname=Caption tag={tagname=P}}");
y -= leading;

id_listitem = p.begin_item("LI", "");
    p.fit_textline("&#x2022;", x1, y, "tag={tagname=Lbl}");
    p.fit_textline("Apples", x2, y, "tag={tagname=LBody}"); y -= leading;
p.end_item(id_listitem);

id_listitem = p.begin_item("LI", "");
    p.fit_textline("&#x2022;", x1, y, "tag={tagname=Lbl}");
    p.fit_textline("Oranges", x2, y, "tag={tagname=LBody}"); y -= leading;
p.end_item(id_listitem);

id_listitem = p.begin_item("LI", "");
    p.fit_textline("&#x2022;", x1, y, "tag={tagname=Lbl}");
    p.fit_textline("Bananas", x2, y, "tag={tagname=LBody}"); y -= leading;
p.end_item(id_listitem);

p.end_item(id_list);
```

Tagging lists with tags in Textflow option lists. The structure in figure 11.6 can alternatively be created with the following tags in Textflow option lists (indented for clarity):

```
<tagbegin={tagname=L ListNumbering=None}>
  <tagbegin={tagname=Caption tag={tagname=P}}>
    The following kinds of fruit are available:\n
  <tagend>
  <tagbegin={tagname=LI}>
    <tagbegin={tagname=Lbl}>&#x2022;<tagend>
    <tagbegin={tagname=LBody}>Apples\n<tagend>
  <tagend>
  <tagbegin={tagname=LI}>
    <tagbegin={tagname=Lbl}>&#x2022;<tagend>
    <tagbegin={tagname=LBody}>Oranges\n<tagend>
  <tagend>
  <tagbegin={tagname=LI}>
    <tagbegin={tagname=Lbl}>&#x2022;<tagend>
    <tagbegin={tagname=LBody}>Bananas\n<tagend>
  <tagend>
<tagend>
```

Cookbook A code sample for creating a Tagged PDF list can be found in the topic `list_pdfua1` in the `pdfua` category of the PDFlib Cookbook.

11.2.4 Creating Contents out of Order

As mentioned in Section 11.1.9, »Print Stream Order and Logical Reading Order«, page 276, it is crucial to create the elements in the structure hierarchy in logical reading order. If the application processes page contents in an order which is different from logical reading order (e.g. always top to bottom regardless of column relationships), several PDFlib features can be used to maintain proper ordering in the structure hierarchy:

- ▶ Create child elements out of order with the *index* option. It modifies the position where a new structure element is inserted within its parent element.
- ▶ Create structure elements out of order with the *parent* option. It modifies the parent element where a new structure element is inserted.
- ▶ Jump back and forth in the structure hierarchy with the `PDF_activate_item()` method. It can be used to add more structure elements or content items to some element in the structure hierarchy.

These methods are discussed in more detail below.

Cookbook Code samples for tagging elements out of order can be found in the topics `tag_out_of_order_pdfua1` and `parallel_columns_pdfua1` in the `pdfua` category of the PDFlib Cookbook.

Creating child elements out of order. In order to create child elements within a structure element out of order you can specify a location in the structure tree with the *index* option of `PDF_begin_item()` or suboptions of the *tag* option of various functions. The following code fragment emits text fragments in reverse order, and corrects the order in the structure tree by inserting each new text fragment as the new first child (*index=0*) of the parent element. Since each new element is inserting as the new first child element of the parent, as a result the logical ordering will be the reverse of the creation order:

```
p.fit_textline("three", x, y, "tag={tagname=P index=0}");
y += leading;
p.fit_textline("two", x, y, "tag={tagname=P index=0}");
y += leading;
p.fit_textline("one", x, y, "tag={tagname=P index=0}");
```

You can query the index of the currently active tag within its parent element with `PDF_get_option()` and the `activeitemindex` or the `activeitemkidcount` keyword, and later return to this position in the structure tree. The following code fragment inserts a new element after the element at the stored index:

```
nextindex = p.get_option("activeitemindex", "") + 1;

...create more elements on the same level...

p.fit_textline(text, x, y, "tag={tagname=P index=" + nextindex + "}");
```

Creating structure elements out of order. In order to create child elements at some other location within the structure tree instead of at the current position, use the `parent` option. It must refer to a structure element which has not yet been closed. Since elements created with abbreviated tagging are created and closed in the same method call they cannot be used as target of the `parent` option. You can query the id of the currently active tag with `PDF_get_option()` and the `activeitemid` keyword, and later return to this position in the structure tree:

```
parent_id = p.get_option("activeitemid", "");
...
p.fit_textline(text, x, y, "tag={tagname=P parent=" + parent_id + "}");
```

For even more flexibility the `parent` and `index` options can be used in combination. Use `PDF_suspend/resume_page()` to interrupt a page, continue on another page, and then go back to add more content to the suspended page.

Activating items for complex layouts. In order to facilitate the creation of structure information for complex non-linear page layouts PDFlib offers a feature called item activation. It can be called to activate a previously created structure element in situations where the developer must keep track of multiple structure branches where each branch could span one or more pages. Typical situations which benefit from this technique are the following:

- ▶ multiple columns on a page
- ▶ inserts which interrupt the main text, such as summaries or other non-linear text items
- ▶ tables and illustrations which are placed between columns.

The method `PDF_activate_item()` allows you to switch back and forth between different branches of the structure tree. The »logical order« approach requires the client application to construct the page contents in logical order even if it would be easier to create it in visual order. In contrast, with item activation the contents can be created in visual order (or any other ordering which is convenient for the application). This technique can also be applied if the content spans multiple pages.

In order to work around problems in Acrobat, no content items should be added immediately after calling `PDF_activate_item()`, but only other structure elements.

Querying the currently active structure element. In order to use the *parent* and *index* options or `PDF_activate_item()` some knowledge about the currently active structure element and its children is required. This status information can be maintained by the application, but can also be queried from PDFlib. The method `PDF_get_option()` with the keywords *activeitemid*, *activeitemindex*, *activeitemkidcount*, *activeitemname*, *activeitem-standardname* delivers the id, index, number of child elements, name, and standard name (if it is role mapped) of the current element.

11.3 Importing Tagged PDF Pages with PDI

Cookbook Code samples for importing pages from Tagged PDF documents can be found in the topics `clone_pdfua` and `merge_and_stamp_pdfua1` in the *pdfua* category of the PDFlib Cookbook.

In Tagged PDF mode pages from a Tagged PDF document are imported along with their structure element tags. We call a page from a Tagged PDF document imported in Tagged PDF mode with *usetags=true* a »tagged page«. This status can be queried with the *tagged* keyword of `PDF_info_pdi_page()`. Importing tagged pages works as described below.

Opening a Tagged PDF document. `PDF_open_pdi_document()` checks whether the imported document is compatible to the current PDF/A-1a/2a/3a or PDF/UA mode and reads the structure tree and role map of the imported document.

If the *usetags* option is *false* the document's structure information is ignored, and no tags can be imported from the document. Note that the default of *usetags* is *false* if `PDF_open_pdi_document()` is called in object scope, i.e. before `PDF_begin_document()`.

Note Attribute classes and class maps are not imported.

Cloning input document language. If most or all pages of a PDF document are imported with PDI it is recommended to clone the document language entry (see »Language specification«, page 274) if it is present in the input document. The document language can be cloned with pCOS and the following code fragment:

```
if (p.pcos_get_string(indoc, "type:/Root/Lang").equals("string"))
{
    inputlang = p.pcos_get_string(indoc, "/Root/Lang");
    optlist += " lang=" + inputlang;
}

p.begin_document(filename, optlist);
```

Opening a tagged page. `PDF_open_pdi_page()` selects the structure elements comprising the imported page's contents and filters the tags present on the page. One or more structure elements on the imported page are selected which form the top of the imported structure sub-tree. If the *usetags* option is *false* the page's structure information is completely ignored.

Entries in the imported document's role map are copied to the output document's role map if the corresponding element is used on the page. Conflicting role map entries (i.e. a custom tag is already mapped to a different standard tag in the generated document's role map or a previously imported document) are ignored. However, pages with conflicting role map entries are rejected in PDF/UA mode, i.e. the call to `PDF_open_pdi_page()` fails.

Importing untagged PDF content. Pages from untagged documents and pages opened with `usetags=false` can also be processed, but they are automatically marked as Artifact.

Importing documents with invalid tag structure. PDFlib applies strict checks for the tag nesting rules imposed by ISO 32000-1 as detailed in »PDF 1.7 Structure Element Nesting Rules«, page 266. These checks can also be applied to imported documents and the tag structure created from imported pages. The nesting rules in imported pages are not checked by default. However, these checks can be enabled with the `checktags` option of `PDF_open_pdi_document()`. If `checktags=strict` all tag nesting rules are enforced in `PDF_open_pdi_page()`. If the structure hierarchy of the imported page violates the nesting rules for structure elements the call to `PDF_open_pdi_page()` fails, and `PDF_get_errmsg()` reports an error similar to the following:

```
Structure type 'Document' cannot contain content items,  
but only other structure elements
```

Since many existing real-world Tagged PDF documents violate the tag nesting rules you can address these problems in imported documents with one of the following methods:

- ▶ Inserting an additional tag on top of the imported structure hierarchy (e.g. with the `tags` option of `PDF_fit_pdi_page()`) is useful to fix common problems where the imported page contains content items immediately under the document root.
- ▶ Other problems cannot be solved by inserting additional tags, e.g. when the structure elements for tables or lists are incomplete. You should consider fixing the input documents if possible. If this is not a viable solution you can set `checktags=none` in `PDF_open_pdi_document()` to import tagged PDF pages with a nonconforming tag structure.
- ▶ If the imported structure itself is correct, but conflicts with the generated new tags of the output document you should try to adjust the new tags appropriately. If this is not feasible you can set `checktags=none` in `PDF_begin_document()` to ignore conflicts in the generated tagging structure. This is not allowed in PDF/UA-1 mode.

Since nonconforming input processed with either variant of `checktags=none` may result in nonconforming PDF output this setting is not recommended.

Querying and checking the tags in an imported page. In some situations it can be difficult to properly integrate imported structure elements in the generated new structure hierarchy. In order to assist in this process, several properties of an imported tagged PDF page can be queried with `PDF_info_pdi_page()` after a page has been successfully opened:

- ▶ The keyword `fittingpossible` reports whether the page can be placed in the current context. If the `tag` option is supplied you can check whether the page can be placed with an additional top-level tag. Only the `tagname` suboption of the `tag` option is evaluated; other suboptions should not be supplied. Using this keyword is recommended in situations where you are not sure about the structure information in imported pages and want to avoid an exception in `PDF_fit_pdi_page()` because of unsuitable imported structure elements. If a page is rejected by the `fittingpossible` test you can try to insert an additional tag via the `tag` option.
- ▶ The keyword `toleveltagcount` reports the number of top-level structure elements since there may be more than one top-level tags. Note that `toleveltagcount` may be 0 in rare cases where the page contents are not covered by any structure element. Such

a page behaves like an untagged page. Like other content items it cannot be placed as child of a grouping element, but requires an additional tag on top of the page.

- ▶ The keyword *toleveltag* reports the top-level structure element(s) of the imported page. This may be useful to determine whether or not additional structure elements can or must be inserted above the imported page structure.
- ▶ The keyword *lang* reports the *lang* attribute of all top-level imported structure element(s). This helps to decide whether a *lang* attribute is required in higher structure elements.

Placing a tagged page. *PDF_fit_pdi_page()* places the imported page on a new page and integrates its structure hierarchy in the generated document's structure tree, using the currently active element as parent for the imported structure tree. An additional structure element can be created with the tag option of *PDF_fit_pdi_page()*. It serves as new parent for the imported structure hierarchy.

If *Alt* or *ActualText* attributes are present in the imported hierarchy they are removed if a conflict with existing attributes higher up in the generated document structure is found (see »Nesting rules for alternate and replacement text«, page 276).

If the page has been opened with *usetags=true* in *PDF_open_pdi_document()* and *PDF_open_pdi_page()* it can be placed at most once in the output document since the imported structure must be integrated at a unique location of the output document's structure hierarchy. If a page contributes content and structure at more than one location it must be opened multiply (i.e. different page handles are placed).

The following strategy is recommended for importing pages with unknown document structure:

- ▶ If the *toleveltagcount* keyword of *PDF_info_pdi_page()* reports a tag count of 0, the page has been imported with *usetags=false*, is empty or contains only Artifacts. Some applications may decide to skip such pages since they don't contribute any relevant content.
- ▶ If the *fittingpossible* keyword of *PDF_info_pdi_page()* returns 1 for this page it can safely be placed with *PDF_fit_pdi_page()*. This test is not required if the page has been loaded with *usetags=false* or is placed as Artifact.
- ▶ Otherwise, an additional tag can be inserted on top of the imported page's structure. The choice of this tag depends on application, especially the type of structure element where the imported page is inserted into the structure tree.
- ▶ If the *fittingpossible* keyword of *PDF_info_pdi_page()* still rejects the page with the additional tag, the application could try another tag or give up the page.

The following code fragment implements the strategy outlined above. It inserts an additional *P* element if the page cannot be placed directly:

```
fittingpossible = true;
additionaltag = "";

toleveltagcount = (int) p.info_pdi_page(page, "toleveltagcount", "");

if (toleveltagcount == 0)
{
    /* The page doesn't contain any structure elements,
     * i.e. it is empty or contains only Artifacts.
     * We add an "Artifact" tag.
     */
    additionaltag = "tag={tagname=Artifact} ";
}
```

```

}
else
/*
* Try to place the page without any additional tag;
* if this doesn't work we insert another tag.
*/
if (p.info_pdi_page(page, "fittingpossible", "") == 0)
{
    additionaltag = "tag={tagname=P} ";

    if (p.info_pdi_page(page, "fittingpossible", additionaltag) == 0)
    {
        fittingpossible = false;
    }
}

if (fittingpossible)
{
    p.fit_pdi_page(page, 0, 0, "adjustpage " + additionaltag);
}
else
{
    System.err.println("Skipping page: " + p.get_errmsg());
}

```

Use cases. The use cases below for importing pages in Tagged PDF mode can be distinguished; Table 11.8 summarizes options and requirements for these use cases:

- ▶ Honor tags: import a page from a Tagged PDF document and copy its tags to the output. The structure elements comprising the page become part of the new structure hierarchy. Optionally an additional tag can be placed on top of the imported hierarchy.
- ▶ Ignore tags: import a page from a Tagged or untagged PDF document and tag the whole page with a new structure element. Existing structure elements are ignored and the contents of the imported page comprise a single element without any internal structure.
- ▶ Place as Artifact: import a page from a Tagged or untagged PDF document, e.g. for a background graphic. The page is automatically tagged as Artifact and existing structure elements are dropped.

11.4 PDF/UA for Universal Accessibility

11.4.1 The PDF/UA-1 Standard

Note General information about the PDF/UA standard can be found on the PDFlib Web site.

Note If you are new to PDF/UA the document »PDF/UA in a Nutshell« published by the PDF Association is recommended (see www.pdfa.org/pdfua-in-a-nutshell/).

PDF/UA-1 improves the universal accessibility of PDF documents by specifying details of Tagged PDF elements and other document aspects. PDF/UA-1 can be regarded as an application of WCAG 2.0 (Web Content Accessibility Guidelines¹) to PDF documents.

1. See www.w3.org/TR/WCAG20/

Table 11.8 Importing tagged and untagged PDF pages in Tagged PDF mode

use case	options	requirements and comments
honor tags	usetags=true in <code>PDF_open_pdi_document()</code> and <code>PDF_open_pdi_page()</code>	<ul style="list-style-type: none"> ▶ parent¹ must not be inline or pseudo element ▶ imported tagged page can be placed only once ▶ top-level element(s) of imported page must meet nesting rules
ignore tags and place page within current element	usetags=false in <code>PDF_open_pdi_document()</code> or <code>PDF_open_pdi_page()</code> and tagname=Span direct=true in <code>PDF_fit_pdi_page()</code>	<ul style="list-style-type: none"> ▶ contents of imported page become part of currently active item ▶ only non-grouping parent elements allowed
place as Artifact	usetags=false in <code>PDF_open_pdi_document()</code> or <code>PDF_open_pdi_page()</code> and tagname=Artifact in <code>PDF_fit_pdi_page()</code>	<ul style="list-style-type: none"> ▶ the currently active element must allow Artifact as child ▶ contents are automatically tagged as Artifact in PDF/UA mode ▶ page can be placed multiply

1. Currently active tag specified with `PDF_begin_item()` or the tag option of `PDF_fit_pdi_page()`

PDF/UA-1 is based on Tagged PDF as defined in PDF 1.7 and ISO 32000-1. It doesn't add any new features to the PDF file format, but mainly makes some accessibility and tagging aspects required which are optional in PDF 1.7. It also clarifies the relationship of different types of structure elements.

PDF/UA-1 as defined in ISO 14289-1. PDF/UA-1 improves the accessibility of PDF documents with the following means:

- ▶ enforcing certain requirements regarding the document structure, i.e. tagging rules;
- ▶ requiring certain pieces of auxiliary information, e.g. metadata and alternate text for graphics;
- ▶ preventing certain PDF elements which defeat the purpose of accessibility.

PDFlib's implementation of PDF/UA-1 is based on the following document:

- ▶ The PDF/UA-1 standard (ISO 14289-1:2014)

Note We recommend the document »Tagged PDF Best Practice Guide: Syntax« published by the PDF association (available at www.pdfa.org/resource/tagged-pdf-best-practice-guide-syntax/)

Creating combined PDF/UA-1 and PDF/A documents. It may be useful to create PDF/UA documents which at the same time conform to PDF/A for archiving (see Section 12.2, »PDF/A for Archiving«, page 305). In fact, if you want to create PDF/A-1a/2a/3a we recommend to adhere to the PDF/UA requirements in order to improve the accessibility of the generated documents. In order to create combined PDF/A and PDF/UA-1 documents supply appropriate values for the `pdfa` and `pdfua` options of `PDF_begin_document()`, e.g.:

```
ret = p.begin_document("combo.pdf", "pdfa=PDF/A-2a pdfua=PDF/UA-1 lang=en");
```

When creating such dual-use documents keep in mind that combo files which adhere to both standards must obey the requirements imposed by both standards. The PDF compatibility level is the minimum of the PDF compatibility of the involved standards; imported PDF documents must adhere to both PDF/UA and PDF/A standards.

We recommend to avoid PDF/A-1a for Tagged PDF output and work with the newer PDF/A-2a or PDF/A-3a standards instead because there are a few conflicts between PDF/UA-1 and PDF/A-1a:

- ▶ PDF/UA-1 requires the page option `taborder=structure` in the presence of annotations. However, the `taborder` option requires PDF 1.5 and cannot be used in PDF/A-1a. As a result, annotations cannot be used in combined PDF/A-1a and PDF/UA-1 documents.
- ▶ The `scope` option is recommended in PDF/UA-1 for *TH* elements, but is only available in PDF 1.5 and cannot be used for PDF/A-1a.
- ▶ Since the *TH*ead and *T*Foot elements require PDF 1.5 they cannot be used for grouping header of footer rows in tables for PDF/A-1a.

Cookbook The `invoice_pdfua1` sample in the *PDFlib Cookbook* demonstrates how to create a combined PDF/UA-1 and PDF/A-2a document.

If PDFlib detects a violation of technical PDF/UA requirements it throws an exception. No PDF output is created in this case. Table 11.9 lists general requirements for creating conforming PDF/UA output.

Table 11.9 General requirements for PDF/UA conformance

item	PDFlib requirements for PDF/UA-1 conformance
general document requirements	<code>PDF_begin/end_document()</code> : the <code>pdfua</code> option must be set to PDF/UA-1 which requires Tagged PDF (the tagged option is automatically set). Operations which require PDF 1.7ext3 or above (e.g. rich media annotations, portfolios) must be avoided. The <code>lang</code> option is required. Option <code>viewerpreferences</code> : only <code>true</code> allowed for suboption <code>displaydoctitle</code> Option <code>permissions</code> : <code>keyword noaccessible</code> not allowed
tagging	All tagging rules (see »PDF 1.7 Structure Element Nesting Rules«, page 266) must be obeyed; the <code>checktags document</code> option must not be set to <code>none</code> .
fonts	The <code>font</code> option <code>embedding</code> must be <code>true</code> . The option <code>unicodemap=false</code> is not allowed. Embedding is also required for Standard Type 1 fonts.
text output and PUA Unicode characters	PUA Unicode characters (e.g. logos and symbols) must have appropriate replacement text specified in the <code>ActualText</code> option of <code>PDF_begin_item()</code> for the enclosing content item or the equivalent tag option of the corresponding output method (see »PUA characters«, page 312).
invisible text	The only exception to the embedding requirement applies to fonts which are exclusively used for invisible text (mainly useful for OCR results), i.e. <code>textrendering=3</code> . This can be controlled with the <code>optimizeinvisible</code> option. If invisible text does not have any rendered equivalent (e.g. a scanned image) it must be marked as <code>Artifact</code> .
layers	Layers can be used, but some options of <code>PDF_define_layer()</code> must be avoided (see PDFlib API Reference).
external content	<code>PDF_begin_template_ext()</code> , <code>PDF_load_graphics()</code> and <code>PDF_open_pdi_page()</code> : the reference option must be avoided.
PDF import	<code>PDF_open_pdi_document()</code> : imported documents must conform to PDF/UA; see Section 11.4.4, »Importing Pages from PDF/UA Documents with PDI«, page 298.
metadata	<code>PDF_set_info()</code> with <code>key=Title</code> or metadata in <code>PDF_begin/end_document()</code> with <code>dc:title</code> in the supplied XMP must be provided with a non-empty value.

11.4.2 Tagging Requirements

Since PDF/UA is based on Tagged PDF the requirements discussed in Section 11.1, »Tagged PDF Basics«, page 257, must be obeyed. However, PDF/UA-1 imposes a number of additional tagging requirements to improve accessibility.

Semantic Requirements. The user must create a document hierarchy and obey the semantic rules listed below. Selecting appropriate structure elements is a crucial component of PDF/UA standard conformance. It is important to understand that the application is responsible for these aspects since PDFlib cannot check them:

- ▶ Tagging must use structure elements which are appropriate for the document structure: if it's a heading, it must be tagged as heading. If it's a table, it must be tagged as a table.
- ▶ Contents which are not meaningful for the document must not be included in the document hierarchy, but instead must be tagged as *Artifact*.
- ▶ Structure elements must be arranged in logical reading order. This can be accomplished most easily by creating the tags in reading order. However, for complex layouts this can also be accomplished with `PDF_activate_item()` (see Section 11.2.4, »Creating Contents out of Order«, page 288).
- ▶ Content must be tagged appropriately if the intended information is not otherwise accessible because of the content's color, format or layout.
- ▶ Text represented in a graphic requires the *Alt* option with an explanation if it doesn't contain text in a natural language (e.g. font or script samples).
- ▶ Image captions must be marked with a *Caption* tag.
- ▶ List elements (*L*) must be created if the content is intended to be read as a list.
- ▶ Headers and footers must be tagged as *Artifact* with `artifacttype=Pageination` and `artifacts subtype=Header` or `Footer`.
- ▶ Only a single *Figure* tag must be created for groups of graphical elements which logically belong together.
- ▶ Footnotes, endnotes, note labels and references to locations within the document must be tagged as *Note* or *Reference* as appropriate.

Note The document »Tagged PDF Best Practice Guide: Syntax« published by the PDF Association provides guidance regarding the correct use of Tagged PDF structure elements for PDF/UA-1.

Tag-specific requirements. Table 11.10 lists requirements for specific tags to achieve PDF/UA-1 conformance. These rules must also be obeyed for custom structure types which are mapped to the listed standard types. For example, if the custom tag *Illustration* is mapped to *Figure* in the `rolemap` option, it is also subject to the conditions for *Figure*.

Table 11.10 Tag-specific requirements for PDF/UA-1 conformance

structure type	PDFlib requirements for PDF/UA-1 conformance
standard structure types	<code>PDF_begin_document()</code> : the <code>rolemap</code> option must not map standard structure types.
<i>Figure</i>	One of the options <code>Alt</code> or <code>ActualText</code> must be supplied.
<i>Formula</i>	The option <code>Alt</code> must be supplied.

Table 11.10 Tag-specific requirements for PDF/UA-1 conformance

structure type	PDFlib requirements for PDF/UA-1 conformance
Table	Table elements must be created for logical tables, but must not be created for tables which are created for layout purposes. Tables formatted by PDFlib can be tagged automatically, see Section 11.2.1, »Automatic Table Tagging«, page 278.
TH	Tables should include headers. The option <code>Scope</code> is recommended for TH elements (it is always created by automatic table tagging if the header option of <code>PDF_fit_table()</code> is used).
L	The structure type L (list) requires the <code>ListNumbering</code> option. If none of its LI (list item) children contains a Lbl (label) element <code>ListNumbering</code> must be <code>None</code> . On the other hand, if <code>ListNumbering=None</code> , but there are visible list labels these should be marked as Artifacts.
Note	The <code>id</code> option is required for footnotes and endnotes tagged as Note (although this doesn't provide any advantage since the <code>id</code> isn't referenced anywhere).

Headings. Appropriate heading tags must be used for all headings. There are two approaches regarding the hierarchic nesting of headings in a PDF document:

- ▶ Strongly structured documents: grouping elements nest as deeply as required to reflect the organization of the content with articles, sections, subsections etc. At each level, the children of the grouping element should consist of a heading *H*, one or more paragraphs *P* for content at that level, and additional grouping elements for nested subsections. Strong structure is typically used in XML documents.
- ▶ Weakly structured documents: the document's structure hierarchy is relatively flat, having only one or two levels of grouping elements, with all the headings, paragraphs, and other BLSEs as their immediate children. The organization of the content is not reflected in the logical structure, but may be expressed by specific heading levels *H1*, *H2*, *H3* etc. Heading tags may not have any descendants. Weak document structure is typically used in HTML.

In PDF/UA-1 mode this distinction must be made explicit with the *structuretype* document option. PDFlib enforces the following rules regarding the use of heading elements depending on the type of document structure.

- ▶ All documents:
 - ▶ The *Title* option should be used in all heading tags to denote document sections (e.g. »Chapter 1«)
 - ▶ Heading elements *H*, *H1*, *H2*, etc. must not have any descendants.
- ▶ Weakly structured documents, i.e. *structuretype=weak* (the default):
 - ▶ Heading sequences must start at *H1* and must not skip any numeric level. For example, the sequence *H1 H3* is not allowed.
 - ▶ Additional heading levels *H7*, *H8* etc. may be used if more than six heading levels are required. Since these are not standard structure types they require an entry in the *rolemap* option. The recommended mapping is *P*.
 - ▶ Unnumbered heading elements *H* must not be used.
- ▶ Strongly structured documents, i.e. *structuretype=strong* (not recommended):
 - ▶ *H* must be used for headings, but there cannot be more than one *H* tag in each node of the structure hierarchy.
 - ▶ Numbered heading elements *H1*, *H2*, etc. must not be used.

11.4.3 Additional Requirements for specific Content Types

Table 11.11 lists PDF/UA-1 requirements and recommendations related to various content types and interactive elements.

Table 11.11 PDF/UA-1 requirements and recommendations for specific content types and interactive elements

Content type	PDFlib requirements for PDF/UA-1 conformance
text	The natural language of all text on a page must be declared; changes of natural language within a sequence of text must also be declared. The natural language can be declared with the lang option of <code>PDF_begin_document()</code> and other means; see »Language specification«, page 274, for details. The language attribute of an element is inherited to all its descendants.
vector graphics and raster images	Raster images and vector graphics must be tagged as Artifact, Figure, or Formula. This affects low-level path construction functions like <code>PDF_rect()</code> etc., path objects with <code>PDF_draw_path()</code> , <code>PDF_fit_image()</code> and similar functions. SVG graphics containing vector graphics or raster images are also affected by this requirement.
imported PDF pages	PDF pages placed with <code>PDF_fit_pdi_page()</code> which contain a graphic should be tagged as Artifact or Figure.
annotations	<code>PDF_begin/end_page_ext()</code> : only the value structure is allowed for the option taborder. <code>PDF_create_annotation()</code> with type=Link: The annotation must be contained in a Link structure element. <ul style="list-style-type: none">▶ The ismap option of <code>PDF_create_action()</code> is not allowed for actions in Link annotations. <code>PDF_create_annotation()</code> with a type different from Link and Popup: <ul style="list-style-type: none">▶ The annotation must be contained in an Annot structure element.▶ The contents option or the tag option with the suboption ActualText is required for visible annotations¹. <code>PDF_fit_pdi_page()</code> : the options useannots and usefields with a value other than none are not allowed if the page has been opened with usestags=false.
form fields	If any form field is present on the page: <code>PDF_begin/end_page_ext()</code> : only the value structure is allowed for option taborder. <code>PDF_create_field()</code> and options fieldname, fieldtype of <code>PDF_add_table_cell()</code> : a Form tag must be created with <code>PDF_create_field()</code> or the tag option. The tooltip option of <code>PDF_create_field()</code> and <code>PDF_create_fieldgroup()</code> is required.
page labels	Page labels created with the labels option in <code>PDF_begin/end_document()</code> and label option in <code>PDF_begin/end_page_ext()</code> should be semantically appropriate.
bookmarks	Generating bookmarks with <code>PDF_create_bookmark()</code> is recommended. The bookmarks should reflect proper reading order and nesting of the content.
attachments	<code>PDF_load_asset()</code> : option description is recommended. Attachments should be accessible in their own right.

1. An annotation is regarded as visible if its rectangle lies at least partially inside the page's CropBox and the display option of `PDF_create_annotation()` is different from hidden and noview.

11.4.4 Importing Pages from PDF/UA Documents with PDI

Additional rules apply when pages from an existing PDF document are imported into a PDF/UA-conforming output document (see Section 8.3, »Importing PDF Pages with PDI«, page 180, for details on PDF import). In order to import tagged pages from existing PDF documents the imported documents and pages must be compatible to the current document according to the following criteria:

- ▶ *PDF_open_pdi_document()*: PDF/UA-1 documents can be imported and the *usetags* option must be *true*. Pages from untagged documents, non-PDF/UA documents or tagged documents loaded with *usetags=false* in *PDF_open_pdi_document()* can also be placed. However, such pages are automatically marked as Artifact.
- ▶ *PDF_open_pdi_page()*: the role map of the imported document must be compatible with the mapping provided by the *rolemap* option of *PDF_begin_document()*. This means that custom structure types must not be mapped to different standard types by the *rolemap* option and the role map of the imported document (or the role map of a previously imported document)
- ▶ *PDF_open_pdi_page()*: the heading structure of the imported page must be compatible with the structure type of the generated document, i.e. if *structuretype=weak* only *H1*, *H2*, etc. (but not *H*) must be used on the page; if *structuretype=strong* only *H* (but not *H1*, *H2*, etc.) must be used on the imported page. Pages with both numbered and unnumbered headings will be rejected.

If PDF/UA-1 conformance is configured in PDFlib and the imported document adheres to the requirements above, the generated output is guaranteed to conform to PDF/UA-1 as well.

Note PDFlib does not validate PDF input documents for PDF/UA conformance nor can it convert arbitrary input PDF documents to PDF/UA.



12 PDF Versions and Standards

Note The PDF/UA standard is discussed in Section 11.4, »PDF/UA for Universal Accessibility«, page 293.

12.1 PDF Versions

12.1.1 PDF Version Overview

PDFlib generates output according to one of the following PDF versions:

- ▶ PDF 1.4 (Acrobat 5, released 2001): mainly useful for creating PDF/A-1
- ▶ PDF 1.5 (Acrobat 6, released 2003)
- ▶ PDF 1.6 (Acrobat 7, released 2005)
- ▶ PDF 1.7 (Acrobat 8, released 2006), technically identical to ISO 32000-1:2008
- ▶ PDF 1.7 Adobe extension level 3 (Acrobat 9, released 2008)
- ▶ PDF 1.7 Adobe extension level 8 (Acrobat X/XI/DC, released 2010/2012/2015-2021)
- ▶ PDF 2.0 according to ISO 32000-2:2017 including dated revision ISO 32000-2:2020

The PDF output version can be controlled with the *compatibility* option in `PDF_begin_document()`. In a particular PDF compatibility mode the PDFlib features for higher levels are not available (see Table 12.1). Trying to use such features results in an exception.

PDF version of documents imported with PDI. In all compatibility modes only PDF documents with a lower or the same PDF version can be imported with PDI. If you must import a PDF with a newer PDF version you must set the *compatibility* option accordingly (see Section 8.3.3, »Document and Page-related Checks«, page 182). As an exception to the *cannot-import-higher-PDF-version* rule, documents according to PDF 1.7 extension level 3 (Acrobat 9) and PDF 1.7 extension level 8 (Acrobat X/XI/DC) can also be imported into PDF 1.7 and PDF 2.0 documents.

12.1.2 PDF 1.x Features

Table 12.1 PDFlib features which require a specific PDF compatibility mode

Feature	PDFlib API methods and options
Features which require PDF 1.7 extension level 8 (Acrobat X/XI/DC) or above	
PDF encryption with AES-256	<code>PDF_begin_document()</code> : AES encryption with 256-bit and a stronger encryption algorithm than in older PDF versions is used if the <code>masterpassword</code> or <code>userpassword</code> option is supplied.
Features which require PDF 1.7 extension level 3 (Acrobat 9) or above	
multimedia	<code>PDF_load_asset()</code> with <code>type=3D/Sound/Video</code> <code>PDF_create_annotation()</code> : option <code>type=RichMedia</code> <code>PDF_create_action()</code> : option <code>type=RichMediaExecute</code>
geospatial PDF	<code>PDF_begin_document()</code> : option <code>viewports</code> <code>PDF_load_image()</code> : option <code>georeference</code>
PDF portfolios with folders	<code>PDF_add_portfolio_folder()</code>

Table 12.1 PDFlib features which require a specific PDF compatibility mode

Feature	PDFlib API methods and options
embed 3D models in PRC format	<code>PDF_load_3ddata()</code> : option type=PRC
Features which require PDF 1.7 = ISO 32000-1 (Acrobat 8) or above	
PDF portfolios	<code>PDF_begin_document()</code> : option portfolio <code>PDF_add_portfolio_file()</code>
Unicode file names for attachments	<code>PDF_begin/end_document()</code> : option attachments, suboption filename
Features which require PDF 1.6 (Acrobat 7) or above	
PDF encryption with AES-128	<code>PDF_begin_document()</code> : options masterpassword or userpassword
NChannel color	<code>PDF_create_devicen()</code> : option subtype=nchannel
user units	<code>PDF_begin/end_document()</code> : option userunit
print scaling	<code>PDF_begin/end_document()</code> : suboption printscaling for viewerpreferences option
document open mode	<code>PDF_begin/end_document()</code> : option openmode=attachments
encrypt file attachments only	<code>PDF_begin/end_document()</code> : option attachmentpassword
attachment description	<code>PDF_begin/end_document()</code> : suboption description for option attachments
embed 3D models in U3D format	<code>PDF_load_3ddata()</code> , <code>PDF_create_3dview()</code> <code>PDF_create_annotation()</code> : option type=3D <code>PDF_create_action()</code> : option type=GoTo3DView
Features which require PDF 1.5 (Acrobat 6) or above	
multimedia	<code>PDF_load_asset()</code> with type=Rendition <code>PDF_create_annotation()</code> : option type=Screen <code>PDF_create_action()</code> : option type=Rendition
DeviceN color space with more than 8 colorants	<code>PDF_create_devicen()</code> : option names
various field options	<code>PDF_create_field()</code> and <code>PDF_create_fieldgroup()</code>
page layout	<code>PDF_begin/end_document()</code> : option pagelayout=twopageleft/right
various annotation options	<code>PDF_create_annotation()</code>
extended permission settings	permissions=plainmetadata in <code>PDF_begin_document()</code> , see Table 3.3
Tagged PDF	various options for <code>PDF_begin_item()</code> ; <code>PDF_begin/end_page_ext()</code> : option taborder
Layers	<code>PDF_define_layer()</code> , <code>PDF_begin_layer()</code> , <code>PDF_end_layer()</code> , <code>PDF_layer_dependency()</code>
JPEG2000 images	imagetype=jpeg2000 in <code>PDF_load_image()</code>
compressed object streams	compressed object streams are automatically generated with compatibility=1.5 or above unless objectstreams=none has been set in <code>PDF_begin_document()</code>

12.1.3 PDF 2.0 Features

PDF 2.0 has been specified in ISO 32000-2:2017 and amended in the dated revision ISO 32000-2:2020. PDF 2.0 and PDF 2.0 incorporates features from the following groups:

- ▶ Acrobat 9 features which are supported in PDFlib with the *compatibility=pdf1.7ext3* document option, e.g. georeferenced PDF, hierarchical Portfolios, and AES-256 encryption; see Table 12.1 for a detailed list.
- ▶ Acrobat X features which are supported in PDFlib with *compatibility=pdf1.7ext8*, particularly AES-256 encryption with a stronger encryption algorithm.
- ▶ Features introduced with the PDF/X, PDF/A-3 or PDF/VT standards such as output intents, associated files and Document Part Hierarchy.
- ▶ Additional features according to Table 12.2.

More details about PDF 2.0 can be found on the PDFlib Web site.

Table 12.2 Features which require PDF 2.0 (ISO 32000-2) or a specific PDF/A, PDF/VT or PDF/X standard

feature	PDFlib API methods and options
tagged PDF	<p>The following structure elements for <i>PDF_begin_item()</i> and the tag option of various methods require PDF 2.0: DocumentFragment, Aside, H7, etc., Title, FENote, Sub, Em, Strong, Artifact (as element of the structure hierarchy)</p> <p>structure destinations: destination options for <i>PDF_add_nameddest()</i> and destination option in <i>PDF_create_action()</i>, <i>PDF_create_annotation()</i>, <i>PDF_create_bookmark()</i> and <i>PDF_begin/end_document()</i> with option <i>structdest</i></p> <p><i>PDF_create_action()</i> with option <i>remotestructdest</i></p> <p>New nesting rules for structure elements have been introduced in PDF 2.0.</p> <p>Several structure element attributes have additional options in PDF 2.0, e.g. <i>artifacttype</i>, <i>artifactsubtype</i>, <i>ListNumbering</i></p> <p>Name spaces for structure elements: options <i>pdftagset</i> and <i>tagset</i> of <i>PDF_end_document()</i> and option <i>namespace</i> of <i>PDF_begin_item()</i></p>
output intents	<p>Document-level output intent ICC profiles require PDF 2.0, PDF/A or PDF/X: <i>PDF_load_iccprofile()</i> with <i>usage=outputintent</i></p> <p>Page-level output intent ICC profiles require PDF 2.0: <i>PDF_begin_page_ext()</i> with option <i>outputintents</i></p>
graphics state	<p><i>PDF_create_gstate()</i> and the <i>gstate</i> option of various methods: options <i>useblackptcomp</i> and <i>halftoneorigin</i></p>
tab order	<p><i>PDF_begin/end_page_ext()</i>: keywords <i>annotations</i> and <i>widgets</i> for option <i>taborder</i></p>
document part hierarchy	<p><i>PDF_begin/end_dpart()</i>: A document part hierarchy requires PDF 2.0 or PDF/VT.</p>
relationship of file attachments	<p>Option relationship for <i>PDF_load_asset()</i> with <i>type=Attachment</i>, for <i>PDF_add_portfolio_file()</i>, and for use as suboption for the <i>attachments</i> option of <i>PDF_begin/end_document()</i> and the attachment option of <i>PDF_create_annotation()</i>: relationship specifications require PDF 2.0 or PDF/A-3.</p>

Table 12.2 Features which require PDF 2.0 (ISO 32000-2) or a specific PDF/A, PDF/VT or PDF/X standard

feature	PDFlib API methods and options
associated files	Option associatedfiles for <code>PDF_end_document()</code> , <code>PDF_begin/end_page_ext()</code> , <code>PDF_begin/end_dpart()</code> , <code>PDF_begin_template_ext()</code> , <code>PDF_load_image()</code> , <code>PDF_open_pdi_page()</code> , <code>PDF_load_graphics()</code> require PDF 2.0 or PDF/A-3. Option associatedfiles for <code>PDF_create_annotation()</code> , <code>PDF_begin_item()</code> and the tag option of various methods requires PDF 2.0 Option structureassociatedfiles for <code>PDF_end_document()</code> requires PDF 2.0.
annotations	<code>PDF_create_annotation()</code> : options blendmode, lang, opacityfill

12.1.4 Deprecated PDF Features

Deprecated features in PDF 2.0. PDF 2.0 deprecates the features listed in Table 12.3 which were available in earlier PDF versions. If these features are used PDFlib emits a warning in PDF 1.x mode and an error in PDF 2.0 mode. We recommend to avoid these features even in PDF 1.x mode.

Table 12.3 Features deprecated in PDF 2.0 which are still available in PDF 1.x mode

PDFlib features deprecated in PDF 2.0	notes
<code>PDF_begin_document()</code> : permissions with keyword noaccessible	Restricting document accessibility is deprecated in PDF 2.0.
<code>PDF_begin/end_document()</code> : viewerpreferences, suboptions printarea, printclip, viewarea, viewclip	Prepress viewer preferences are deprecated in PDF 2.0.
<code>PDF_begin_page_ext()</code> and <code>PDF_end_page_ext()</code> : separationinfo	Separation dictionaries are deprecated in PDF 2.0.
<code>PDF_create_annotation()</code> : type=Movie	Use type=RichMedia.
<code>PDF_create_action()</code> : type=Movie	Use type=RichMediaExecute.
<code>PDF_create_gstate()</code> : blendmode accepts only a single keyword, but not a list with multiple values	Blendmode arrays are deprecated in PDF 2.0.

Features outside of ISO 32000. The features listed in Table 12.4 were never part of ISO 32000-1. They work only in Acrobat, but not in third-party viewers. If any of these features is used, PDFlib emits a warning in PDF 1.x mode and an error in PDF 2.0 mode. We recommend to avoid these features even in PDF 1.x mode since they are not part of standard PDF.

Table 12.4 Features outside of ISO 32000 which are still available in PDF 1.x mode

features outside of ISO 32000	PDFlib API methods and options
barcode fields	<code>PDF_create_field()</code> and <code>PDF_create_fieldgroup()</code> : barcode
search index	<code>PDF_begin/end_document()</code> : search
editable watermarks	<code>PDF_begin_template_ext()</code> : watermark

12.2 PDF/A for Archiving

12.2.1 The PDF/A Standards

Note General information about the PDF/A standard can be found on the PDFlib Web site.

The PDF/A formats specified in the ISO 19005 standard series provide a consistent and robust subset of PDF which can safely be archived over a long period of time, or used for reliable data exchange in enterprise and government environments.

PDF/A in the PDF Association. PDFlib GmbH is a founding member of the PDF Association which hosts the PDF/A Technical Working Group (TWG) as one of its activities. The aim of this industry organization is »to promote Open Standards-based electronic document implementations using PDF technology through education, expertise and shared experience for stakeholders worldwide«. For more information visit the PDF Association's Web site at www.pdfa.org.



PDF/A-1a:2005 and PDF/A-1b:2005 according to ISO 19005-1. PDF/A-1 is based on PDF 1.4 and imposes restrictions on the use of color, fonts, annotations, and other elements. There are two flavors of PDF/A-1:

- ▶ ISO 19005-1 Level B conformance (PDF/A-1b) ensures that the visual appearance of a document is preservable over the long term. Simply put, PDF/A-1b ensures that the document will look the same when it is processed some time in the future.
- ▶ ISO 19005-1 Level A conformance (PDF/A-1a) is based on level B, but adds properties which are known from Tagged PDF: it requires structure information and reliable text semantics in order to preserve the document's logical structure and natural reading order. PDF/A-1a not only ensures that the document will look the same when it is processed in the future, but also that its contents can be reliably interpreted and will be accessible to physically impaired users.

PDF/A-1 support in PDFlib is based on the following documents:

- ▶ The PDF/A-1 standard (ISO 19005-1:2005)
- ▶ Technical Corrigendum 1 (ISO 19005-1:2005/Cor.1:2007)
- ▶ Technical Corrigendum 2 (ISO 19005-1:2005/Cor.2:2011)
- ▶ TechNote 0010 »Clarifications of ISO 19005« published by the PDF Association.

When PDF/A-1 (without any conformance level) is mentioned, both conformance levels PDF/A-1a and PDF/A-1b are meant.

PDF/A-2a, PDF/A-2b, and PDF/A-2u according to ISO 19005-2. The PDF/A-2 standard flavors are based on ISO 32000-1 (i.e. PDF 1.7), which means that they support more features than PDF/A-1. Unlike PDF/A-1 the newer PDF/A-2 standard allows transparency, layers, JPEG 2000 compression, PDF/A file attachments, PDF packages and other PDF features. PDF/A-2 defines the following flavors:

- ▶ ISO 19005-2 Level B conformance (PDF/A-2b) ensures the visual appearance of a document.

- ▶ ISO 19005-2 Level A conformance (PDF/A-2a) adds reliable Unicode text semantics and Tagged PDF with structure information. The tags make sure that PDF/A-2a documents are fully accessible.
- ▶ ISO 19005-2 Level U conformance (PDF/A-2u) sits in between PDF/A-2a and PDF/A-2b in that it requires reliable Unicode text semantics, but not structure information. PDF/A-2u guarantees that the pages can faithfully be reproduced and that the text can be extracted and searched.

PDF/A-2 support in PDFlib is based on the following documents:

- ▶ The PDF/A-2 standard (ISO 19005-2:2011)
- ▶ TechNote 0010 »Clarifications of ISO 19005« published by the PDF Association.

When PDF/A-2 (without any conformance level) is mentioned, all three conformance levels PDF/A-2a, PDF/A-2b, and PDF/A-2u are meant.

PDF/A-3a, PDF/A-3b, and PDF/A-3u as defined in ISO 19005-3. PDF/A-3 is similar to PDF/A-2 with the following differences:

- ▶ While PDF/A-2 allows only file attachments which conform to PDF/A-1 or PDF/A-2, PDF/A-3 allows arbitrary file types as attachments.
- ▶ Attached files are associated with the whole document, a page, or some other element of the document. The relationship between the file attachment and the corresponding part of the document must be specified explicitly, e.g. source, alternative, or supplemental data.

PDF/A-3 support in PDFlib is based on the following document:

- ▶ The PDF/A-3 standard (ISO 19005-3:2012)
- ▶ TechNote 0010 »Clarifications of ISO 19005« published by the PDF Association.

When PDF/A-3 (without any conformance level) is mentioned, all three conformance levels PDF/A-3a, PDF/A-3b, and PDF/A-3u are meant.

The ZUGFeRD standard for electronic invoices is an important application based on PDF/A-3. It embeds a machine-readable XML version of the invoice in a human-readable document which conforms to PDF/A-3. More information on ZUGFeRD can be found on the PDFlib Website.

12.2.2 General Requirements

Cookbook Code samples for generating PDF/A can be found in the pdfa category of the PDFlib Cookbook.

If the PDFlib client program obeys to the rules documented in this chapter, valid PDF/A output is guaranteed. If PDFlib detects a violation of PDF/A rules it will throw an exception which must be handled by the application. No PDF output is created in this case. Table 12.5 lists general requirements for creating conforming PDF/A output.

Creating combined PDF/A and PDF/UA-1 documents. A PDF/A document can at the same time conform to PDF/UA-1. In fact, if you want to create PDF/A-1a/2a/3a we recommend to adhere to the PDF/UA requirements in order to improve the accessibility of the generated documents. See »Creating combined PDF/UA-1 and PDF/A documents«, page 294, for details and restrictions.

Table 12.5 General requirements for PDF/A conformance levels A, B, and U

item	PDFlib requirements for PDF/A conformance (all conformance levels)
PDF/A conformance level and PDF compatibility	PDF_begin_document() : the pdfa option must be set to the required PDF/A conformance level, e.g. pdfa=PDF/A-2b PDF/A-1: Operations that require PDF 1.5 or above (e.g. layers) must be avoided. PDF/A-2/3: Operations that require PDF 1.7ext3 or above (e.g. PDF Portfolios) must be avoided.
fonts	The font option embedding must be true. The option unicodemap=false is not allowed. Embedding is also required for the Standard Type 1 fonts. The only exception to the embedding requirement applies to fonts which are exclusively used for invisible text (mainly useful for OCR results). This can be controlled with the optimizeinvisible option.
page sizes	PDF_begin/end_page_ext() : there are no strict page size limits in PDF/A. However, it is recommended to keep the page size (width and height, and all box entries) in the range 3...14400 points (508 cm) in PDF/A-1, or 3...14400 user units in PDF/A-2/3.
layers	PDF/A-1: PDF_define_layer() and PDF_set_layer_dependency() must be avoided. PDF/A-2/3: layers can be used, but some options of PDF_define_layer() must be avoided.
security	PDF_begin_document() : the userpassword, masterpassword, attachmentpassword and permissions options must be avoided.
external content	PDF_begin_template_ext() , PDF_load_graphics() and PDF_open_pdi_page() : the reference option must be avoided. PDF_load_asset() : the external option must be avoided.
file size	The file size of the generated PDF document must not exceed 2 GB, and the number of PDF objects must be smaller than 8,388,607. See Section 3.1.6, »Maximum Size of PDF Documents and other Limits«, page 56, for more details about these limits.
PDF import	PDF_open_pdi_document() is restricted unless infomode=true; see Section 12.2.7, »Importing PDF/A Documents with PDI«, page 312.

Creating combined PDF/A and PDF/X documents. A PDF/A document can at the same time conform to PDF/X-3 or PDF/X-4, but not to PDF/X-4p or PDF/X-5n. In order to create such a combo file supply appropriate values for the *pdfa* and *pdfx* options of **PDF_begin_document()**, e.g.:

```
ret = p.begin_document("combo.pdf", "pdfa=PDF/A-2b pdfx=PDF/X-4");
```

12.2.3 Color and Image Requirements

PDF/A guarantees faithful color reproduction by requiring device-independent color specifications. Color spaces may come from the following sources:

- ▶ images loaded directly with **PDF_load_image()** and **PDF_fill_imageblock()**, and indirectly via **PDF_load_graphics()**
- ▶ explicit color specifications with **PDF_set_graphics_option()** or **PDF_setcolor()**
- ▶ color specifications through option lists, e.g. in Textflows
- ▶ blending color spaces for transparency groups: **PDF_begin/end_page_ext()**, **PDF_begin_template_ext()**, and **PDF_load_graphics()**: option *transparencypgroup* with suboption *colorspace*
- ▶ alternate color space of spot or DeviceN colors
- ▶ annotations and form fields may specify colors for borders, background and contents.

Table 12.6 lists PDF/A requirements for color processing which must be obeyed in all of the operations listed above.

Table 12.6 Color and image requirements for PDF/A conformance levels A, B, and U

item	PDFlib requirements for PDF/A (all conformance levels)
output condition (output intent)	<i>PDF_load_iccprofile()</i> with usage=outputintent or <i>PDF_process_pdi()</i> with action=copyoutputintent must be called immediately after <i>PDF_begin_document()</i> if any of the device-dependent colors spaces Gray, RGB, or CMYK is used in the document and no suitable default color space is present for the page.
grayscale color	Grayscale color can only be used if a grayscale, RGB, or CMYK output intent is present or the defaultgray option has been set.
RGB color	RGB color for page content can only be used if an RGB output intent is present or the defaultrgb option has been set. RGB color can always be used in form fields and annotations.
CMYK color	CMYK color can only be used if a CMYK output intent is present or the defaultcmk option has been set.
Separation (spot) and DeviceN color	<ul style="list-style-type: none"> ▶ The alternate color space must conform to the rules above. ▶ PDF/A-2/3: <i>PDF_makespotcolor()</i> must be called before <i>PDF_create_devicen()</i> for all custom spot colors in the DeviceN color space.
transparency and overprinting	<p>PDF/A-1: Transparency must be avoided; this affects the following API features:</p> <ul style="list-style-type: none"> ▶ <i>PDF_load_image()</i>: the masked option must be avoided unless the mask refers to a 1-bit image. ▶ <i>PDF_load_image()</i>: images with implicit transparency (alpha channel) are not allowed; they must be loaded with the ignoremask option. ▶ <i>PDF_load_graphics()</i>: SVG graphics containing transparent elements must be avoided. ▶ <i>PDF_create_gstate()</i>: the opacityfill and opacitystroke options must be avoided unless they have a value of 1; if blendmode is used it must be Normal; if softmask is used it must be none. ▶ <i>PDF_create_annotation()</i>: the opacity option must be avoided. <p>PDF/A-2/3: Transparency is allowed, but the following rule must be obeyed in <i>PDF_create_gstate()</i>: overprintmode=1 is not allowed if the current color space is an ICCBased CMYK color space and overprintfill or overprintstroke is true.</p>
transparency groups	<p><i>PDF_begin/end_page_ext()</i>, <i>PDF_begin_template_ext()</i>, <i>PDF_open_pdi_page()</i> and <i>PDF_load_graphics()</i>: the option transparencygroup is restricted as follows:</p> <ul style="list-style-type: none"> ▶ PDF/A-1: Option transparencygroup is not allowed. ▶ PDF/A-2/3: The suboption colorspace of the transparencygroup option must meet the requirements stated above for grayscale, RGB, and CMYK color. For <i>PDF_open_pdi_page()</i> and <i>PDF_load_graphics()</i> transparencygroup=auto is enforced. ▶ PDF/A-2/3: The suboption colorspace of the transparencygroup option must meet the requirements stated above for Grayscale, RGB and CMYK color. The option transparencygroup=none and the suboption colorspace=none are not allowed for <i>PDF_begin/end_page_ext()</i>.
images and templates	<p><i>PDF_load_image()</i>: the interpolate=true option must be avoided.</p> <p>PDF/A-2/3: JPEG 2000 images must meet certain conditions, see »JPEG 2000 images«, page 161, for details.</p>

Output intents. The output condition defines the intended target device, which is important for consistent color rendering. Unlike PDF/X, which always requires an output intent, the use of an output intent ICC profile is optional in PDF/A. An output intent is only required if device-dependent colors, e.g. RGB, are used in the document. If only device-independent colors, e.g. ICC-based colors, are used in the document no output intent is required. While PDF/X supports only printer ICC profiles as output intents, in PDF/A also monitor profiles are allowed. This makes it possible to use the widely used sRGB profile as output intent. The output intent can be specified with an ICC profile as follows:

```
icc = p.load_iccprofile("sRGB", "usage=outputintent");
```

As an alternative to loading an ICC profile, the output intent can be copied from an imported PDF/A document (see »Copying the PDF/A output intent from an imported document«, page 313). The output intent of the generated output document must be set exactly once. It should be set immediately after `PDF_begin_document()`.

Color strategies for creating PDF/A. The summary of color strategies in Table 12.7 may be helpful for planning PDF/A applications. The easiest approach which works in many situations is to use the sRGB output intent ICC profile since it supports grayscale and RGB color. In addition, sRGB is known to PDFlib internally and thus doesn't require any external profile data or configuration.

In order to create black text without the need for any output intent profile the CIE Lab color space can be used. The Lab color value (0, 0, 0) specifies pure black in a device-independent manner, and conforms to PDF/A without any output intent profile (unlike DeviceGray, which requires an output intent profile). PDFlib initializes the current color to black at the beginning of each page. Depending on whether or not an ICC output intent was specified, it will use the DeviceGray or Lab color space for black. Use the following call to manually set Lab black color:

```
p.set_graphics_option("fillcolor={lab 0 0 0}");
```

Table 12.7 PDF/A color strategies for conformance levels A, B, and U

output intent	color spaces which can be used in the document					
	CIE Lab	ICCBased	separation and DeviceN	Grayscale ¹	RGB ^{1, 2}	CMYK ¹
none	yes	yes	yes	–	–	–
grayscale ICC profile	yes	yes	yes	yes	–	–
RGB ICC profile	yes	yes	yes	yes	yes	–
CMYK ICC profile	yes	yes	yes	yes	–	yes

1. Device color space without any ICC profile or default color space for the page, pattern or template
 2. RGB color is always allowed for form fields and annotations.

In addition to the color spaces listed in Table 12.7, custom spot colors can be used subject to the corresponding alternate color space. Since PDFlib uses CIE Lab as the alternate color space for the builtin HKS and PANTONE spot colors, these can always be used with the PDF/A standard. For custom spot colors the alternate color space must be chosen so that it is compatible with the output intent. For all custom spot colors used in a DeviceN color space `PDF_makespotcolor()` must be called.

12.2.4 Requirements for Interactive Features

Table 12.8 lists all operations which are restricted when generating PDF/A-conforming output. Calling one of the prohibited functions in PDF/A mode triggers an exception.

Table 12.8 Requirements for interactive features for all PDF/A conformance levels

item	PDFlib requirements for PDF/A (all conformance levels)
annotations	<p>PDF/A-1: <code>PDF_create_annotation()</code> is subject to the following restrictions:</p> <ul style="list-style-type: none"> ▶ <code>PDF_create_annotation()</code>: <code>type=FileAttachment</code> and <code>Highlight</code> are not allowed. The following types are only allowed if an RGB output intent has been specified: <code>Freetext</code>, <code>Text</code>, <code>Stamp</code>, <code>Underline</code>, <code>Squiggly</code>, <code>Strikeout</code>, <code>Polygon</code>, <code>Polyline</code>, <code>Line</code>, <code>Square</code>, <code>Circle</code>, <code>Ink</code>. ▶ The zoom and rotate options for Text annotations must not be set to true. ▶ The option <code>richtext</code> is not allowed. ▶ Options <code>annotcolor</code>, <code>fillcolor</code>, <code>interiorcolor</code>: grayscale colors are only allowed with an output intent (any type), ▶ The opacity option must not be used. <p>PDF/A-2/3: <code>PDF_create_annotation()</code> is subject to the following restrictions:</p> <ul style="list-style-type: none"> ▶ <code>PDF_create_annotation()</code>: with <code>type=3D</code> or <code>Screen</code> is not allowed. ▶ Options <code>annotcolor</code>, <code>fillcolor</code>, <code>interiorcolor</code>: grayscale colors are only allowed with an output intent (any type); CMYK colors are only allowed with a CMYK output intent ▶ The zoom and rotate options for Text annotations must not be set to true. ▶ The option <code>richtext</code> is not allowed.
attachments	<p>PDF/A-1: <code>PDF_begin/end_document()</code>: the attachments option must be avoided.</p> <p>PDF/A-2: <code>PDF_begin/end_document()</code>: the attachments option must refer to PDF/A-1 or PDF/A-2 documents.</p> <p>PDF/A-3: Arbitrary file types can be attached with the <code>associatedfiles</code> option, but the attachments option must be avoided. The following conditions must be obeyed:</p> <ul style="list-style-type: none"> ▶ Attachments can be associated with various parts of the document with the <code>associatedfiles</code> option of <code>PDF_end_document()</code>, <code>PDF_begin/end_page_ext()</code>, <code>PDF_begin/end_dpart()</code>, <code>PDF_begin_template_ext()</code>, <code>PDF_create_annotation()</code>, <code>PDF_load_image()</code>, <code>PDF_open_pdi_page()</code>, <code>PDF_load_graphics()</code>. Each attachment must be associated with exactly one part of the document, i.e. each asset handle created with <code>PDF_load_asset()</code> must be supplied to exactly one <code>associatedfiles</code> option. ▶ The <code>mimetype</code> and <code>relationship</code> suboptions are required. ▶ The <code>description</code> suboption is recommended. ▶ The <code>external=true</code> suboption must be avoided.
actions and JavaScript	<p><code>PDF_create_action()</code>: actions with <code>type=Hide</code>, <code>Launch</code>, <code>Movie</code>, <code>ResetForm</code>, <code>ImportData</code>, <code>JavaScript</code> must be avoided; for <code>type=name</code> only <code>NextPage</code>, <code>PrevPage</code>, <code>FirstPage</code>, and <code>LastPage</code> are allowed.</p> <p><code>PDF_begin/end_document()</code>: the option <code>action</code> can only be used with the trigger event <code>open</code>.</p> <p><code>PDF_begin/end_page_ext()</code>: option <code>action</code> must be avoided.</p>

Table 12.8 Requirements for interactive features for all PDF/A conformance levels

item	PDFlib requirements for PDF/A (all conformance levels)
form fields	<p>PDF <code>create_field/fieldgroup()</code> are subject to the following restrictions:</p> <ul style="list-style-type: none"> ▶ All used fonts must be embedded. ▶ Options <code>backgroundcolor</code>, <code>bordercolor</code>, <code>fillcolor</code>, and <code>strokecolor</code>: RGB colors are always allowed, grayscale colors are only allowed with an output intent (any type), and CMYK colors are only allowed with a CMYK output intent (see also Table 12.6). ▶ The action option is not allowed.

12.2.5 Additional PDF/A Requirements for Level U Conformance

Most standard requirements for PDF/A-2u and PDF/A-3u are met automatically by PDFlib. When generating documents with Level U conformance only the operation listed in Table 12.9 is restricted. In other words, if your application already creates PDF/A-2b or PDF/A-3b and you observe the restriction in Table 12.9 the generated documents can also be declared as PDF/A-2u or PDF/A-3u, respectively.

Table 12.9 Additional restriction for PDF/A conformance level U

item	PDFlib method and option requirements for PDF/A-2u/3u conformance
fonts	The font option <code>unicodemap=false</code> must be avoided.

12.2.6 Additional PDF/A Requirements for Level A Conformance

When creating PDF/A-1a, PDF/A-2a and PDF/A-3a all Tagged PDF requirements according to Section 11.1, »Tagged PDF Basics«, page 257, must be met. Table 12.10 lists required and recommended operations for generating output according to Level A. In addition to the general Tagged PDF rules we recommend to obey the PDF/UA requirements to improve the accessibility of the generated documents; see Section 11.4, »PDF/UA for Universal Accessibility«, page 293, for details.

The user is responsible for creating correct structure information. A document which contains all of its text in a single structure element is technically correct PDF/A, but violates the goal of faithful semantic reproduction.

Table 12.10 Additional requirements for PDF/A conformance level A

item	PDFlib requirements for PDF/A-1a/2a/3a conformance
fonts	The font option <code>unicodemap=false</code> must be avoided.
Tagged PDF	All requirements for Tagged PDF must be met (see Section 11.1, »Tagged PDF Basics«, page 257). The structure hierarchy of the document should reflect the logical structure of the document as accurately as possible.
word boundaries	Words must be separated by space characters (U+0020). The <code>autospace</code> option can be used to simplify this task.
text output and PUA Unicode characters	PDF/A-2a/3a: PUA Unicode characters (e.g. logos and symbols) must have appropriate replacement text specified in the <code>ActualText</code> option of PDF <code>begin_item()</code> for the enclosing content item or the equivalent tag option of the corresponding output method (see below for details).
annotations	PDF <code>create_annotation()</code> : the <code>contents</code> option is recommended for annotations which do not display any text.

PUA characters. PDF/A-2a and PDF/A-3a include additional requirements for characters with a Unicode value in the Private Use Area or PUA, i.e. mainly the range U+E000 - U+F8FF (see »BMP and PUA«, page 95, for details). PUA characters are usually decorative and symbolic glyphs, or custom glyphs such as a company logo. PDF/A-2a/3a require that PUA characters are accompanied by an *ActualText* attribute which contains a textual representation of the character. The *ActualText* may be assigned to an individual PUA character or to a longer sequence of characters which includes a PUA character. It is recommended to supply the *ActualText* option with a *Span* inline-level element.

You can use `PDF_info_font()` to check the Unicode value of a particular code for a specific font (see Section 6.5.2, »Font-specific Encoding, Unicode, and Glyph Name Queries«, page 132):

```
uv = (int) p.info_font(font, "unicode", "code=" + c);
```

If the resulting Unicode value *uv* falls into the PUA it requires an *ActualText* attribute. The following code fragment assumes a font called *PDFlibLogo* which includes a graphical representation of the PDFlib corporate logo. When placing the logo on a page a *Span* element with a suitable *ActualText* suboption containing the text *PDFlib Logo* is supplied in the *tag* option:

```
p.fit_textline(text, 50, 700,  
    "fontname=PDFlibLogo fontsize=24 " +  
    "tag={tagname=Span ActualText={PDFlib Logo}}");
```

If you don't have any information about the glyph and therefore no suitable *ActualText* is readily available you may use the name of the glyph in the font. It can be determined as follows:

```
gn_idx = (int) p.info_font(font, "glyphname", "code=" + c);  
glyphname = p.get_option(gn_idx, "");
```

The glyph name may be used in *ActualText*, probably in combination with a fixed phrase. For example, code *ox1A* in the *Wingdings* font contains an image of a computer keyboard with the glyph name *keyboard*. This glyph maps to U+F037, i.e. a PUA value. Using the *ActualText symbol for keyboard* may make sense for this symbol. It should be noted that programmatically constructing *ActualText* must be considered a makeshift solution. Human-selected text is always preferable to machine-generated *ActualText*.

12.2.7 Importing PDF/A Documents with PDI

Additional rules apply when pages from an existing PDF document are to be imported into a PDF/A-conforming output document (see Section 8.3, »Importing PDF Pages with PDI«, page 180, for details on PDF import). All imported documents must conform to a PDF/A conformance level which is compatible to the current PDF/A mode according to Table 12.11.

Note *PDFlib does not validate PDF input documents for PDF/A conformance, nor can it convert arbitrary input PDF documents to PDF/A.*

If a certain PDF/A conformance level is configured in PDFlib and the imported documents adhere to a compatible level, the generated output is guaranteed to conform to the selected PDF/A conformance level. Documents which are incompatible to the current PDF/A level will be rejected in `PDF_open_pdi_document()`.

Table 12.11 Compatible PDF/A input levels for various PDF/A output levels

PDF/A output level	PDF/A level of the imported document				
	PDF/A-1a:2005	PDF/A-1b:2005	PDF/A-2a, PDF/A-3a	PDF/A-2b, PDF/A-3b	PDF/A-2u, PDF/A-3u
PDF/A-1a:2005	allowed	–	–	–	–
PDF/A-1b:2005	allowed	allowed	–	–	–
PDF/A-2a, PDF/A-3a	allowed	–	allowed	–	–
PDF/A-2b, PDF/A-3b	allowed	allowed	allowed	allowed	allowed
PDF/A-2u, PDF/A-3u	allowed	–	allowed	–	allowed

Cookbook A full code sample can be found in the Cookbook topic `pdfa/c/clone_pdfa`.

If one or more PDF/A documents are imported, they must all have been prepared for a compatible output condition according to Table 12.12. The output intents in all imported documents must be identical or compatible; it is the user’s responsibility to make sure that this condition is met.

Table 12.12 Output intent compatibility when importing PDF/A documents (all conformance levels)

output intent of generated document	output intent of imported document			
	none	grayscale	RGB	CMYK
none	yes	–	–	–
grayscale ICC profile	yes	yes ¹	–	–
RGB ICC profile	yes	–	yes ¹	–
CMYK ICC profile	yes	–	–	yes ¹

1. The output intent of the imported document and the output intent of the generated document must be identical.

While PDFlib can correct certain items, it is not intended to work as a PDF/A validator or to enforce PDF/A conformance for imported documents. For example, PDFlib will not embed fonts which are missing from imported PDF pages.

If you want to combine imported pages such that the resulting PDF output document conforms to the same PDF/A conformance level and output condition as the input document(s), you can query the PDF/A status of the imported PDF as follows:

```
pdfalevel = p.pcos_get_string(doc, "pdfa");
```

This statement retrieves a string designating the PDF/A conformance level of the imported document if it conforms to a PDF/A level, or *none* otherwise. The returned string can be used to set the PDF/A conformance level of the output document appropriately, using the `pdfa` option in `PDF_begin_document()`.

Copying the PDF/A output intent from an imported document. In addition to querying the PDF/A conformance level you can also copy the PDF/A output intent from an imported document. Since PDF/A documents do not necessarily contain any output intent you must first use pCOS to check for the existence of an output intent before attempting to copy it.

Cookbook A full code sample can be found in the *Cookbook* topic `pdfa/clone_pdfa`.

This can be used as an alternative to setting the output intent via `PDF_load_iccprofile()`, and will copy the imported document's output intent to the generated output document. Copying the output intent works for imported PDF/A and PDF/X documents.

12.2.8 XMP Metadata for PDF/A

PDF/A relies on the XMP format for embedding metadata in PDF documents. PDF/A supports two kinds of XMP metadata: a set of well-known metadata schemas called predefined schemas which are taken from the underlying version of the XMP specification, and custom extension schemas. PDFlib automatically creates the required PDF/A conformance entries in the XMP as well as several common entries (e.g. *CreationDate*).

Document-level XMP. XMP document metadata can be supplied with the *metadata* option of `PDF_begin_document()`, `PDF_end_document()`, or both. In PDF/A mode PDFlib verifies whether user-supplied XMP document metadata conforms to the PDF/A requirements. XMP metadata from imported PDF documents can be fetched from the input PDF via the pCOS path `/Root/Metadata`.

Cookbook A full code sample can be found in the *Cookbook* topic `interchange/import_xmp_from_pdf`.

Component-level XMP. In addition to the document as a whole, XMP metadata can also be attached to other components in a PDF document such as pages or images. While there are no PDF/A-1 requirements for component-level metadata, PDF/A-2 and PDF/A-3 mandate that custom properties in component-level XMP are also described by an extension schema description similar to document-level XMP.

Component-level XMP metadata can be supplied with the *metadata* option of `PDF_begin/end_page_ext()`, `PDF_load_image()` and other functions.

Predefined XMP schemas. Use of XMP for document metadata in PDF/A is based on the following specifications:

- ▶ PDF/A-1: XMP 2004 specification
- ▶ PDF/A-2 and PDF/A-3: XMP 2005

The schemas described in the respective XMP specification are called predefined schemas, and are listed in Table 12.13 along with their namespace URI and the preferred namespace prefix. Only properties of predefined schemas can be used with PDF/A unless an extension schema description is present (see below). The full list of properties in the predefined XMP 2004 schemas for PDF/A-1 is available in TechNote 0008 from the PDF/A Competence Center of the PDF Association. PDF/A-2/3 add predefined schemas from XMP 2005, but the additional schemas are related to images and dynamic media, and are therefore unlikely to be useful for document metadata.

Table 12.13 Predefined XMP schemas for PDF/A (see XMP 2004 and XMP 2005 for details)

Schema name and description	namespace URI	preferred namespace prefix
XMP 2004 schemas for use in PDF/A-1, PDF/A-2, and PDF/A-3		
Adobe PDF schema	<code>http://ns.adobe.com/pdf/1.3/</code>	pdf
Dublin Core schema	<code>http://purl.org/dc/elements/1.1/</code>	dc

Table 12.13 Predefined XMP schemas for PDF/A (see XMP 2004 and XMP 2005 for details)

Schema name and description	namespace URI	preferred namespace prefix
EXIF schema for EXIF-specific properties	http://ns.adobe.com/exif/1.0/	exif
EXIF schema for TIFF properties	http://ns.adobe.com/tiff/1.0/	tiff
Photoshop schema	http://ns.adobe.com/photoshop/1.0/	photoshop
XMP Basic Job Ticket schema	http://ns.adobe.com/xap/1.0/obj	xmpBJ
XMP Basic schema	http://ns.adobe.com/xap/1.0/	xmp
XMP Media Management schema	http://ns.adobe.com/xap/1.0/mm/	xmpMM
XMP Paged-Text schema	http://ns.adobe.com/xap/1.0/t/pg/	xmpTPg
XMP Rights Management schema	http://ns.adobe.com/xap/1.0/rights/	xmpRights
Additional XMP 2005 schemas for use in PDF/A-2 and PDF/A-3		
Camera Raw Schema	http://ns.adobe.com/camera-rawsettings/1.0/	crs
EXIF Schema for Additional EXIF Properties	http://ns.adobe.com/exif/1.0/aux/	aux
XMP Dynamic Media Schema	http://ns.adobe.com/xmp/1.0/DynamicMedia/	xmpDM

XMP extension schema descriptions. If your metadata requirements are not covered by the predefined schemas you can define an XMP extension schema. PDF/A describes an extension mechanism which must be used when custom schemas are to be embedded in a document. Table 12.14 summarizes the schemas which must be used for describing one or more extension schemas and their properties, along with their namespace URI and the required namespace prefix. Note that the namespace prefixes are required (unlike the preferred namespace prefixes for predefined schemas).

If custom XMP properties for component-level XMP (e.g. on page level) are used, the corresponding extension schema description can be supplied together with the custom XMP properties in the same method (e.g. `PDF_begin_page_ext()`). As an alternative, the extension schema description for component-level XMP can be supplied with the document-level XMP in `PDF_begin_document()`.

More details of constructing XMP extension schema descriptions and an example are available in TechNote 0009 from the PDF/A Competence Center.

Cookbook Full code and XMP samples can be found in the Cookbook topics `pdfa/xmp_extension_schema_pdfa2b` and `pdfa/pdfa_extension_schema_with_type_pdfa2b`.

Table 12.14 PDF/A extension schema container schema and auxiliary schemas

Schema name and description	namespace URI ¹	required namespace prefix
PDF/A extension schema container schema: container for all embedded extension schema descriptions	http://www.aiim.org/pdfa/ns/extension/	pdfaExtension
PDF/A schema value type: describes a single extension schema with an arbitrary number of properties	http://www.aiim.org/pdfa/ns/schema#	pdfaSchema

Table 12.14 PDF/A extension schema container schema and auxiliary schemas

Schema name and description	namespace URI ¹	required namespace prefix
PDF/A property value type: describes a single property	http://www.aiim.org/pdfa/ns/property#	pdfaProperty
PDF/A ValueType value type: describes a custom value type used in extension schema properties; only required if types beyond the XMP 2004 list of types are used.	http://www.aiim.org/pdfa/ns/type#	pdfaType
PDF/A field type schema: describes a field in a structured type	http://www.aiim.org/pdfa/ns/field#	pdfaField

1. The namespace URIs are incorrectly listed in ISO 19005-1, and have been corrected in Technical Corrigendum 1.

12.3 PDF/X for Print Production

12.3.1 The PDF/X Family of Standards

Note If you are new to PDF/X the document »PDF/X in a Nutshell« published by the PDF Association is recommended as an introduction (see www.pdfa.org/pdfx-in-a-nutshell/).

The PDF/X formats specified in ISO 15930 define a consistent and robust subset of PDF which can be used to exchange data suitable for commercial printing. PDFlib can generate output and process input conforming to the PDF/X flavors described below. Table 12.15 summarizes the main differences between the PDF/X flavors.

Table 12.15 Comparison of PDF/X flavors

	PDF/X-3	PDF/X-4	PDF/X-4p	PDF/X-5n
PDF version	PDF 1.4	PDF 1.6	PDF 1.6	PDF 1.6
CMYK and spot color	yes	yes	yes	yes
device-independent color (ICCBased and Lab)	yes	yes	yes	yes
transparency and layers	–	yes	yes	yes
externally referenced output intent ICC profile	–	–	yes	yes
n-colorant output intent ICC profile	–	–	–	yes

PDF/X-1a:2003 as defined in ISO 15930-4. This standard is based on PDF 1.4, with some features (e.g. transparency) prohibited. PDF/X-1a supports only CMYK and spot colors, but not RGB or ICC-based colors. For this reason PDF/X-1a is deprecated.

PDF/X-3:2003 as defined in ISO 15930-6. This standard is based on PDF 1.4 and supports workflows based on device-independent ICC-based color in addition to grayscale, CMYK, and spot colors. Output devices can be monochrome, RGB, or CMYK. Some PDF 1.4 features, especially transparency, are prohibited. Because of the very old PDF version and the exclusion of transparency PDF/X-3 is deprecated.

PDF/X-4 as defined in ISO 15930-7. This standard is based on PDF 1.6. In PDF/X-4 transparency and layers are allowed, but some other PDF 1.6 features are still prohibited. The variant PDF/X-4p allows output intent ICC profiles to be kept external from the PDF document to save space. PDFlib implements 15930-7:2010.

PDF/X-5n as defined in ISO 15930-8. This standard is targeted at »partial exchange« which requires prior discussion between producer and receiver of a file. PDF/X-5n is based on PDF 1.6 and supports external output intent ICC profiles for n-colorant print characterizations, also called xCLR profiles.

PDFlib implements ISO 15930-8:2010 including Corrigendum 1.

12.3.2 General Requirements

Cookbook Code samples for generating PDF/X can be found in the pdfx category of the PDFlib Cookbook.

If the PDFlib client program obeys to the rules documented in this chapter, valid PDF/X output is guaranteed. If PDFlib detects a violation of PDF/X rules it throws an exception and no PDF output is created. Table 12.16 lists general requirements for creating conforming PDF/X output.

Table 12.16 General requirements for PDF/X conformance

item	PDFlib requirements for PDF/X compatibility
PDF/X conformance level and PDF compatibility	<p>PDF_begin_document(): the pdfx option must be set to the required PDF/X conformance level, e.g. pdfx=PDF/X-4.</p> <p>PDF/X-3: Operations that require PDF 1.5 or above must be avoided.</p> <p>PDF/X-4/5n: Operations that require PDF 1.7 or above must be avoided.</p>
fonts	<p>The font option embedding must be true. Embedding is also required for the PDF core fonts. The only exception to the embedding requirement applies to fonts which are exclusively used for invisible text (mainly useful for OCR results). This can be controlled with the optimizeinvisible option (which is not allowed in PDF/X-3).</p>
page sizes	<p>PDF_begin/end_page_ext(): The page boxes, which are settable via the cropbox, bleedbox, trimbox, and artbox options, must satisfy the following requirements:</p> <ul style="list-style-type: none"> ▶ The TrimBox or ArtBox must be set, but not both of these box entries. If both TrimBox and ArtBox are missing PDFlib will take the CropBox (if present) as the TrimBox, and the MediaBox if the CropBox is also missing. ▶ The BleedBox, if present, must fully contain the ArtBox and TrimBox. ▶ The CropBox, if present, must fully contain the ArtBox and TrimBox.
layers	<p>PDF/X-3: Layers must be avoided.</p> <p>PDF/X-4/5n: Layers can be used, but some options of PDF_define_layer() must be avoided.</p>
document info fields and XMP metadata	<p>The Creator and Title info fields must be set to a non-empty value with PDF_set_info() or (in PDF/X-4/5n) with the xmp:CreatorTool and dc:title XMP properties in the metadata option of PDF_begin/end_document()</p> <p>Values other than True or False for the Trapped info field in PDF_set_info() or the corresponding XMP property pdf:Trapped in the metadata option of PDF_begin/end_document() must be avoided.</p>
security	<p>PDF_begin_document(): the userpassword, masterpassword, and permissions options must be avoided.</p>

Table 12.16 General requirements for PDF/X conformance

item	PDFlib requirements for PDF/X compatibility
file size	PDF/X-4/5n: The file size of the generated PDF document must not exceed 2 GB, and the number of PDF objects must be smaller than 8.388.607. See Section 3.1.6, »Maximum Size of PDF Documents and other Limits«, page 56, for more details about these limits.
PDF import	PDF <code>open_pdi_document()</code> is restricted unless <code>infomode=true</code> ; see Section 12.3.6, »Importing PDF/X Documents with PDI«, page 322.

12.3.3 Output Intent and Color Requirements

PDF/X guarantees faithful color reproduction by requiring device-independent color specifications. Color spaces may come from the following sources:

- ▶ images loaded directly with `PDF_load_image()` and `PDF_fill_imageblock()`, and indirectly via `PDF_load_graphics()`
- ▶ explicit color specifications with `PDF_set_graphics_option()` or `PDF_setcolor()`
- ▶ color specifications through option lists, e.g. in Textflows
- ▶ alternate color space of spot or DeviceN colors
- ▶ blending color spaces for transparency groups: `PDF_begin/end_page_ext()`, `PDF_begin_template_ext()`, and `PDF_load_graphics()`: option `transparencygroup` with suboption `colorspace`

The operations above are subject to various restrictions when creating PDF/X output as detailed below.

Output intents. The output intent defines the intended target device or printing condition. It is useful for proofing (simulating a particular printing device on another device), characterizing device-specific colors and as blending color space for transparency calculations. The PDF/X output intent is usually described by a Grayscale, RGB or CMYK ICC printer profile with the exception of PDF/X-5n which supports n-colorant ICC profiles. The details vary among PDF/X flavors:

- ▶ PDF/X-3/4: the output intent can be specified by embedding an ICC profile for the target device or printing condition.
- ▶ PDF/X-4p and PDF/X-5n: by referencing an external ICC profile for the output intent (the *p* in the name of the standard means that an external profile is referenced). A strong reference to the ICC file is created which is protected by a checksum. Although a referenced ICC profile is not necessarily embedded in the PDF output, it must nevertheless be available at PDF creation time to create a strong reference. The `urls` option must be provided with one or more valid URLs where the ICC profile can be found:

```
if (p.load_iccprofile("CGATS TR 001",
"usage=outputintent urls={http://www.color.org}") == -1)
{
    /* Error */
}
```

- ▶ PDF/X-5n: in order to make the output intent ICC profile available to the recipient of the PDF document, the ICC profile can be attached to the PDF document via the `embedprofile` option. This attaches the ICC profile to the PDF document similar to a file attachment. If the profile is embedded this way, the `urls` option is not required.

Table 12.17 compares output intent handling in all supported PDF/X flavors. The ICC profile for the output intent should be specified by the print provider based on the intended printing process and paper. If the printing process is unknown or the print provider doesn't specify any output intent you can find recommendations for selecting a suitable ICC profile at www.pdflib.com.

Table 12.17 Output intent handling for various PDF/X flavors

	PDF/X-3	PDF/X-4	PDF/X-4p	PDF/X-5n
output intent color space	Grayscale, RGB or CMYK	Grayscale, RGB or CMYK ICC profile		n-colorant ICC profile (xCLR)
output intent ICC profile embedding	always embedded	always embedded	external reference, but ICC profile can optionally be attached	
other restrictions	–	CMYK output intent profile cannot be used for other purposes	output intent profile cannot be used for other purposes	

Color requirements for PDF/X-3 and PDF/X-4/4p. PDF/X-3/4 support device-independent workflows with ICC profiles. Table 12.18 lists requirements for these standards. Since PDF/X-5n targets different workflows its requirements are listed separately in the next section.

A special rule applies to PDF/X-4: a CMYK output intent profile (i.e. loaded with *usage=outputintent*) cannot be used for an ICCBased color space (i.e. loaded with *usage=iccbased*) in the same document. This requirement is mandated by the PDF/X standard and applies only to CMYK profiles, but not to Grayscale or RGB profiles. If you run into this conflict you can simply omit the ICC profile and use plain CMYK color instead since the output intent CMYK profile is applied anyway. In order to avoid unnecessary error messages Pdflib ignores a CMYK ICC profile embedded in an image or SVG graphics if it is identical to the PDF/X output intent.

The restriction still applies to imported documents: if an imported page uses the same CMYK ICC profile as the generated document's output intent it is rejected by `PDF_open_pdi_page()`.

Table 12.18 Color requirements for PDF/X-3 and PDF/X-4/4p conformance

item	Pdflib requirements for PDF/X-3 PDF/X-4/4p compatibility
output condition (output intent)	<code>PDF_load_iccprofile()</code> with <i>usage=outputintent</i> or <code>PDF_process_pdi()</code> with <i>action=copyoutputintent</i> must be called immediately after <code>PDF_begin_document()</code> . A Grayscale, RGB or CMYK output intent must be provided.
grayscale color	Grayscale color can only be used if a Grayscale or CMYK output intent is present or the default-gray option has been set.
RGB color	RGB color can only be used if an RGB output intent is present or the defaultrgb option has been set.
CMYK color	CMYK color can only be used if a CMYK output intent is present or the defaultcmky option has been set.
Separation (spot) and DeviceN color	<ul style="list-style-type: none"> ▶ PDF/X-3/4: The alternate color space must conform to the rules above. ▶ PDF/X-4: <code>PDF_makespotcolor()</code> must be called before <code>PDF_create_devicen()</code> for all custom spot colors in the DeviceN color space. ▶ PDF/X-4: the <i>colorspace</i> suboption of the <i>process</i> option of <code>PDF_create_devicen()</code> must match the PDF/X output intent.

Output intent and color requirements for PDF/X-5n. The output intent ICC profile for PDF/X-5n must describe an n-colorant printing device. Such profiles contain the color space designation *xCLR*, where *x* is a hexadecimal digit in the range 2 to F, i.e. *2CLR...FCLR* for 2...15 colorants. *xCLR* profiles are not supported in plain PDF, but can only be used as PDF/X-5n output intent. As PDF doesn't support such profiles directly, they cannot be embedded as output intent, but must be referenced as external file. As an external object, however, it can be embedded as file attachment to enhance availability.

xCLR profiles contain a colorant table with the names and color values of named colorants. With respect to the allowed color spaces in a PDF/X-5n document it is relevant whether the *xCLR* output intent profile contains the colorant *Black*, or possibly all of the CMYK colorants *Cyan*, *Magenta*, *Yellow*, *Black*. For example, a *7CLR* profile for a printing process with CMYK, Orange, Green and Violet colorants (CMYKOGV) satisfies both conditions.

Table 12.19 lists requirements for PDF/X-5n.

Table 12.19 Color requirements for PDF/X-5n conformance

item	PDFlib requirements for PDF/X-5n compatibility
output condition (output intent)	<i>PDF_load_iccprofile()</i> with <code>usage=outputintent</code> or <i>PDF_process_pdi()</i> with <code>action=copyout-intent</code> must be called immediately after <i>PDF_begin_document()</i> . An n-colorant (<i>xCLR</i>) profile must be provided.
grayscale color	Grayscale color can only be used if the output intent contains the colorant <i>Black</i> or the default-gray option has been set.
RGB color	RGB color can only be used if the defaultrgb option has been set.
CMYK color	CMYK color can only be used if the output intent contains all of the colorants <i>Cyan</i> , <i>Magenta</i> , <i>Yellow</i> and <i>Black</i> or the defaultcmky option has been set.
Separation (spot) and DeviceN color	<ul style="list-style-type: none"> ▶ The alternate color space must conform to the rules above. ▶ <i>PDF_makespotcolor()</i> must be called before <i>PDF_create_devicen()</i> for all custom spot colors in the DeviceN color space. Spot colors in the colorant list of the output intent are exempted from this requirement. ▶ The options <code>subtype=nchannel</code> and process of <i>PDF_create_devicen()</i> are not allowed.

Choosing a suitable PDF/X output intent. The PDF/X output intent is usually selected as a result of discussions between you and the print service provider who takes care of print production. If your printer cannot provide any information regarding the choice of output intent, you can find suitable ICC output intent profiles on the Internet.

Color strategies for PDF/X-3/4/5. Regarding color handling the following two strategies can be implemented, possibly in combination:

- ▶ Device-independent color: regardless of the type of output intent ICC profile, device-independent color spaces can be used, i.e. ICC-based or CIELab. The Lab color value (*o, o, o*) specifies pure black in a device-independent manner. Use the following call to manually set Lab black color:

```
p.set_graphics_option("fillcolor={lab 0 0 0}");
```

- ▶ Device-dependent color: device-specific grayscale, RGB, or CMYK color can be used. While grayscale color can be used with any kind of output intent, RGB or CMYK colors can only be used with a matching output intent.

Spot and DeviceN colors can be used subject to the corresponding alternate color space. Since PDFlib uses CIELab as the alternate color space for the builtin HKS and Pantone spot colors, these can always be used with PDF/X-3/4/5. For custom spot colors the alternate color space must be chosen so that it is compatible with the output intent. For all custom spot colors used in a DeviceN color space `PDF_makespotcolor()` must be called.

The summary of color strategies in Table 12.20 may be helpful for planning PDF/X applications.

Table 12.20 PDF/X-3/4/5 color strategies

output intent ICC profile	color spaces which can be used in the document					
	CIELab	ICCBased	Separation, DeviceN	Grayscale ¹	RGB ¹	CMYK ¹
grayscale	yes	yes	yes	yes	–	–
RGB ²	yes	yes	yes	yes	yes	–
CMYK	yes	yes	yes	yes	–	yes
xCLR for PDF/X-5n	yes	yes	yes	yes	–	–

1. Device color space without any ICC profile or default color space for the page, pattern or template

2. The PDF/X standard requires a printer profile; monitor profiles, e.g. sRGB, can not be used. RGB printer profiles are very rare.

12.3.4 Image and Transparency Requirements

Raster images must obey the color-related requirements discussed in Section 12.3.3, »Output Intent and Color Requirements«, page 318. Table 12.21 lists additional image-related requirements for PDF/X conformance.

Table 12.21 image requirements for PDF/X conformance

item	PDFlib requirements for PDF/X compatibility
images	<ul style="list-style-type: none"> ▶ PDF/X-3: JBIG2 images must be avoided. ▶ PDF/X-4/5: JPEG 2000 images must satisfy certain conditions, see »JPEG 2000 images«, page 161, for details.
transparency	<p>PDF/X-3: Transparency must be avoided; this affects the following API features:</p> <ul style="list-style-type: none"> ▶ <code>PDF_load_image()</code>: the masked option must be avoided unless the mask refers to an image with 1-bit depth. ▶ <code>PDF_load_image()</code>: images with implicit transparency (alpha channel) are not allowed; they must be loaded with the ignoremask option. ▶ <code>PDF_load_graphics()</code>: SVG graphics containing transparent elements must be avoided. ▶ <code>PDF_create_gstate()</code>: the opacityfill and opacitystroke options must be avoided unless they have a value of 1; if blendmode is used it must be Normal; if softmask is used it must be none. <p>PDF/X-4/5n: Transparent images and graphics can be used.</p>

Table 12.21 image requirements for PDF/X conformance

item	PDFlib requirements for PDF/X compatibility
transparency groups	<p>PDF_begin/end_page_ext(), PDF_begin_template_ext(), PDF_open_pdi_page() and PDF_load_graphics(): the option <code>transparencygroup</code> is restricted as follows:</p> <ul style="list-style-type: none"> ▶ PDF/X-3: The option <code>transparencygroup</code> is not allowed. ▶ PDF/X-4: The suboption <code>colorspace</code> of the <code>transparencygroup</code> option must meet the requirements in Table 12.18 for Grayscale, RGB and CMYK color. For PDF_open_pdi_page() and PDF_load_graphics() <code>transparencygroup=auto</code> is enforced. ▶ PDF/X-5n: The suboption <code>colorspace</code> of the <code>transparencygroup</code> option must meet the requirements stated in Table 12.19 for Grayscale, RGB and CMYK color. The option <code>transparencygroup=none</code> and suboption <code>colorspace=none</code> are not allowed for PDF_begin/end_page_ext().

12.3.5 Requirements for interactive Features

Table 12.22 lists all operations which are restricted when generating PDF/X-conforming output. Calling one of the prohibited functions in PDF/X mode triggers an exception.

Table 12.22 PDF/X requirements for interactive features

item	PDFlib method and option requirements for PDF/X compatibility
annotations and form fields	<p>PDF_create_annotation(), PDF_create_field(): annotations inside the <code>BleedBox</code> (or <code>TrimBox/ArtBox</code> if no <code>BleedBox</code> is present) must be avoided. For the following types the annotation rectangle must be located outside the <code>BleedBox</code> at least one half the line width: <code>Square</code>, <code>Circle</code>, <code>Ink</code>, <code>PolyLine</code>, <code>Polygon</code>. The options <code>calloutline</code>, <code>cloudy</code> and <code>createrichtext</code> of PDF_create_annotation() is not allowed.</p> <p>PDF/X-3: PDF_create_annotation() with <code>type=Caret</code>, <code>FileAttachment</code> and <code>Highlight</code> is not allowed.</p> <p>PDF/X-4/5: PDF_create_annotation() with <code>type=Caret</code>, <code>Highlight</code> and <code>Screen</code> is not allowed.</p> <p>PDF_fit_pdi_page() with options <code>useannots/usefields</code>: annotations and form fields are only imported if they are located inside the <code>BleedBox</code> (or <code>TrimBox/ArtBox</code> if no <code>BleedBox</code> is present) of the generated page. The option <code>clipannotations=false</code> is required in this case.</p>
file attachments	PDF/X-3: PDF_begin/end_document() : option <code>attachments</code> must be avoided.
actions and JavaScript	PDF_create_action() : all actions including JavaScript must be avoided.
viewer preferences / view and print areas	PDF_begin/end_document() : if the <code>viewarea</code> , <code>viewclip</code> , <code>printarea</code> , and <code>printclip</code> suboptions for the <code>viewerpreferences</code> option are used values other than <code>media</code> or <code>bleed</code> are not allowed.

12.3.6 Importing PDF/X Documents with PDI

Additional rules apply when pages from an existing PDF document are imported into a PDF/X output document (see Section 8.3, »Importing PDF Pages with PDI«, page 180, for details). All imported documents must conform to a PDF/X level which is compatible to the current PDF/X mode according to Table 12.23. For all allowed combinations with PDF/X-4/5 output the following additional rule must be observed: if an imported page uses the same CMYK ICC profile as the generated document's output intent, it is rejected by **PDF_open_pdi_page()** since this would violate the PDF/X-4/5 standard.

If a certain PDF/X conformance level is configured in PDFlib and the imported documents adhere to a compatible level, the generated output is guaranteed to conform to

the selected PDF/X conformance level. Documents which are incompatible to the current PDF/X level are rejected in `PDF_open_pdi_document()`.

Table 12.23 Compatible PDF/X input levels for various PDF/X output levels

PDF/X output level	PDF/X level of the imported document				
	PDF/X-1a, PDF/X-3	PDF/X-4	PDF/X-4p	PDF/X-5n	
PDF/X-3	allowed				
PDF/X-4	allowed	allowed	allowed		
PDF/X-4p	allowed	allowed	allowed ¹		
PDF/X-5n					allowed ²

1. `PDF_process_pdi()` with `action=copyoutputintent` copies the reference to the external output intent ICC profile.

2. The imported document must use the same output intent (identical checksum).

If multiple PDF/X documents are imported, they must all have been prepared for the same output condition.

While PDFlib can correct certain items, it is not intended to work as a PDF/X validator or to enforce PDF/X compatibility for imported documents. For example, PDFlib will not embed fonts which are missing from imported PDF pages, and does not apply any color correction to imported pages.

If you want to combine imported pages such that the resulting PDF output document conforms to the same PDF/X conformance level and output condition as the input document(s), you can query the PDF/X status of the imported PDF as follows:

```
pdfxlevel = p.pcos_get_string(doc, "pdfx");
```

This statement retrieves a string designating the PDF/X conformance level of the imported document if it conforms to an ISO PDF/X level, or *none* otherwise. The returned string can be used to set the PDF/X conformance level of the output document appropriately, using the `pdfx` option in `PDF_begin_document()`.

Copying the PDF/X output intent from an imported document. In addition to querying the PDF/X conformance level you can also copy the output intent from an imported document.

Cookbook A full code sample can be found in the Cookbook topic `pdfx/clone_pdfx`.

This can be used as an alternative to setting the output intent via `PDF_load_iccprofile()`, and will copy the imported document's output intent to the generated output document. Copying the output intent works for imported PDF/A and PDF/X documents.

12.4 PDF/VT for Variable and Transactional Printing

12.4.1 The PDF/VT Standard

Note General information about the PDF/VT standard can be found on the PDFlib Web site. The document »PDF/VT Application notes« published by the PDF Association is also recommended (see www.pdfa.org/resource/pdfvt-application-notes/).

The PDF/VT standard has been published as ISO 16612-2. It is »designed to enable variable document printing (VDP) in a variety of environments«. PDF/VT documents contain the final content elements and associated metadata, but not any variables or templates. PDF/VT is based on PDF/X-4 and supports PDF 1.6 features including transparency and layers. In addition to the requirements of PDF/X the PDF/VT standard adds supplementary features to PDF to meet the requirements of high-volume personalized printing. PDF/VT enables high-performance rendering (*RIPping*) of digital print files by adding efficient resource management to PDF.

PDF/VT-1, based on PDF 1.6, is designed for single-file exchange and is based on PDF/X-4. All resources required for rendering a PDF document are contained in a single PDF/VT-1 file.

12.4.2 PDF/VT Concepts

This section provides an overview of the technical concepts on which PDF/VT is based.

Document part hierarchy. The document part (DPart) hierarchy specifies the sequence and relationship of documents or parts of documents in a PDF/VT file. In a common scenario the PDF/VT file contains sub-documents for many recipients and each document part corresponds to the pages comprising the sub-document for a single recipient. In addition to assigning pages to recipients, the document part hierarchy can also reflect more complex structures. For example, the recipients may be grouped according to the ZIP code in their address, the ZIP codes may be organized according to the state, the states according to the countries, etc. This kind of document organization creates a tree-like structure which includes all pages in the document. The elements of this tree are called DPart nodes, where each inner node contains other DPart nodes and each leaf node specifies one or more pages for a recipient.

The document part hierarchy in a PDF/VT file can be used for accessing pages, alternatively to other methods such as access by page number or by page label. The DPart hierarchy is required in PDF/VT files. The optional record level value selects the level in the DPart hierarchy which corresponds to a record for an individual recipient. This is relevant for the scope hints (see below).

Document Part Metadata (DPM). Each node in the document part hierarchy from the root node down to the leaves in the document tree may contain Document Part Metadata (DPM). It can be used to communicate information about a recipient's document and its parts. In particular, properties which are relevant for production (e.g. number of copies of a document part) or information about the recipient (e.g. the corresponding ZIP code) can be encoded in DPM.

JDF (or other) production metadata is not required in PDF/VT, but adds substantial value in JDF-enabled workflows.

Optimizations for recurring graphical content. Print elements are often re-used on multiple pages, e.g. a company logo or product image which appears on all pages of a mailing. Optimized processing of recurring graphical content is an important strategy for improving file size and processing speed of print files. PDF has always supported XObjects as a means for optimizing the file size by including the required data for a print element only once in the file, and allowing references to this data from multiple pages (or multiple instances on the same page). XObjects may contain a raster image or arbitrary text and vector graphics contents. While XObjects in PDF aim at optimizing the overall size of a document, PDF up to now did not include any means for optimized rendering of repeated page contents. There is nothing in PDF which could tell the consuming software that, say, an image on a particular page will appear again on another page later in the same document or maybe in the next print job. PDF/VT extends the existing concept of XObjects in PDF and adds the following means for optimizing print performance:

- ▶ *Unique identification:* XObjects can be assigned an identifier (called *GTS_XID*) which is unique across all documents. This identifier can be used for caching implementations which need to identify equivalent XObjects. In simple terms, the graphic which has already been processed for job 1 and is found to be re-used in job 2 does not have to be ripped again, but the rendered results can be taken from the cache.
- ▶ *Scope hints and environment context:* XObjects may contain information (called *GTS_Scope*) regarding the range of pages or documents where the graphical content is re-used. This way XObjects can carry information about the useful lifetime of their rendered results in the cache: will the content be re-used only for the current recipient, re-used anywhere else in the same file or file stream, or not at all. If an environment context (called *GTS_Env*) is provided, the XObject may specify global use, i.e. it will be re-used in more than one PDF/VT instance. There are no restrictions regarding the environment context string. For example, a customer or job name can be used for identifying the environment.
- ▶ *Encapsulation hints:* XObject caching algorithms in RIPs must take into account the interaction of an XObject with the calling context and existing print elements on the same page (or other pages, e.g. when imposing multiple pages on the same sheet). For example, if an XObject does not specify the color or line width, but varies its appearance based on the color and line width in effect when it is referenced, caching of the rendered result is ineffective due to the varying appearance. A similar situation arises if the XObject contains transparent elements so that the existing background must be blended with the XObject. In order to facilitate caching of XObjects in the RIP, PDF/VT introduces the concept of Encapsulated XObjects which can be marked as such using the *GTS_Encapsulated* key. Encapsulated XObjects must satisfy certain rules which facilitate caching in the RIP.

All of these entries are optional: PDF/VT does not require any of the optimizations for recurring graphical content, but using them offers significant print performance advantages with PDF/VT capable RIPs.

12.4.3 Summary of Rules for generating PDF/VT

Cookbook A code sample for generating PDF/VT-1 can be found in the `pdfvt` category of the *PDFlib Cookbook*.

Creating PDF/VT with PDFlib is achieved by the following means:

- ▶ Since PDF/VT is based on PDF/X, all requirements for the underlying PDF/X conformance level must be met. The `pdfvt` and `pdfx` document options must be used with suitable values.
- ▶ Document parts must be specified with the `DPart` API. Each node in the document part hierarchy may optionally carry DPM metadata. The record level in the `DPart` tree can be specified.
- ▶ Scope hints for recurring graphical content should be provided.
- ▶ If the document contains transparency additional effort should be spent in order to satisfy the conditions for encapsulated XObjects.
- ▶ Additional rules apply when importing pages from existing PDF documents (see Section 12.4.7, »Importing PDF/X and PDF/VT Documents with PDI«, page 330).

Table 12.24 summarizes required and optional operations for generating PDF/VT-conforming output in addition to the corresponding PDF/X requirements listed in Section 12.3.2, »General Requirements«, page 317. Specific aspects are discussed in more detail in subsequent sections below.

Table 12.24 Required and optional operations for PDF/VT conformance

item	PDFlib requirements for PDF/VT conformance
conformance level	<p>The <code>pdfvt</code> option in <code>PDF_begin_document()</code> must be set to the required PDF/VT conformance level, e.g. <code>pdfvt=PDF/VT-1</code>.</p> <p>PDF/VT-1: The option <code>pdfx=PDF/X-4</code> is automatically set; supplying any other value for the <code>pdfx</code> option is an error.</p>
document part hierarchy	<p>The document part hierarchy must be specified. This involves the following operations:</p> <ul style="list-style-type: none">▶ <code>PDF_begin_document()</code>: The <code>nodenamelist</code> option must specify names for all levels of the <code>DPart</code> tree. The <code>recordlevel</code> option can be used to specify the level of the <code>DPart</code> tree which corresponds to recipient records.▶ <code>PDF_begin_dpart()</code> and <code>PDF_end_dpart()</code>: The document part hierarchy must be constructed.▶ Document parts may optionally carry document part metadata (DPM). DPM can be generated by constructing a DPM dictionary with the POCA interface.
scope hints for recurring graphical content	<p><code>PDF_load_image()</code>, <code>PDF_open_pdi_page()</code>, <code>PDF_begin_template_ext()</code> and <code>templateoptions</code> option of <code>PDF_load_graphics()</code>: The <code>scope</code> suboption should be supplied to the <code>pdfvt</code> option with one of the values <code>unknown</code>, <code>singleuse</code>, <code>record</code>, <code>file</code>, <code>stream</code>, or <code>global</code> to provide usage hints for recurring images, templates, and imported PDF pages.</p> <p>For <code>scope=stream</code> and <code>scope=global</code> the suboption <code>environment</code> must specify a PDF/VT environment context, i.e. an identifier that can be used by a conforming PDF/VT reader to provide a management interface for managing related XObjects. For example, the customer name or job name could be used to identify the environment.</p>
encapsulated XObjects	<p><code>PDF_begin_document()</code>: If the document doesn't contain any transparent elements the option <code>usestransparency=false</code> should be supplied so that PDFlib can encapsulate all Form XObjects without a transparency group.</p> <p><code>PDF_load_image()</code>: Images should use the <code>renderingintent</code> or <code>mask</code> option to allow the XObject to be marked as encapsulated (see Section 12.4.6, »Encapsulated XObjects«, page 329).</p>

Creating combined PDF/VT and PDF/A documents. It may be useful to create PDF/VT print documents which at the same time conform to PDF/A for archiving (see Section 12.2, »PDF/A for Archiving«, page 305, for more information on PDF/A). Keep the following in mind when creating such dual-use documents:

- ▶ The following document options can be used to create PDF/VT and PDF/A at the same time:

```
ret = p.begin_document("combo.pdf", "pdfx=PDF/X-4 pdfvt=PDF/VT-1 pdfa=PDF/A-2b");
```

- ▶ The restrictions stated in this manual for PDF/VT and PDF/A must be obeyed, and only features which are allowed in both standards can be used.
- ▶ Imported PDF documents must adhere to both the PDF/X and PDF/A standards in order to be acceptable.

12.4.4 Document Part Hierarchy and Document Part Metadata (DPM)

Cookbook Code samples for generating PDF/VT with a document part hierarchy can be found in the pdfvt category of the PDFlib Cookbook.

Creating the document part hierarchy. All pages in the document must be organized in the document part hierarchy. In simple cases, e.g. invoices, this tree structure consists of two levels, the *root* level and the *recipient* level. The pages in the document comprise a linear arrangement of recipient records, where each record contains one or more pages. The structure of the document part hierarchy must be specified in the document option list, e.g.

```
if (p.begin_document(outfile,  
    "pdfvt=PDF/VT-1 nodenamelist={root recipient} recordlevel=1") == -1)
```

More complex documents may require a deeper document hierarchy, e.g. customized brochures consisting of front, body, and back parts which are maintained in the *docpart* level:

```
if (p.begin_document(outfile,  
    "pdfvt=PDF/VT-1 nodenamelist={root recipient docpart} recordlevel=1") == -1)
```

The *recordlevel* option specifies a zero-based index into the *nodenamelist* where the recipient or record level can be found, which is relevant for the *scope=record* option.

After defining the structure of the hierarchy the pages must be grouped in document hierarchy nodes. This is achieved with *PDF_begin/end_dpart()*:

```
/* Create root node of the DPart hierarchy */  
p.begin_dpart("");  
    p.begin_dpart("");          /* Create a new node at the recipient level */  
        p.begin_page_ext(...); /* Create one or more pages */  
        ...  
        p.end_page_ext(...);  
    p.end_dpart("");          /* Close recipient node */  
  
    p.begin_dpart("");          /* Create next node at the recipient level */  
        p.begin_page_ext(...); /* Create one or more pages */  
        ...  
        p.end_page_ext(...);  
    p.end_dpart("");          /* Close recipient node */
```

```

/* Close root node */
p.end_dpart("");

```

The calls to `PDF_begin/end_dpart()` must create a tree structure according to the `node-namelist` option, i.e. the maximum nesting level must correspond to the number of entries in the `nodenamelist` array.

The `groups` option of `PDF_begin_document()` and the `group` and `pagenumber` options of `PDF_begin/end_page_ext()` are not allowed in PDF/VT mode since they would interfere with the DPart hierarchy.

Creating Document Part Metadata (DPM). For each node in the document part hierarchy metadata information can be supplied which applies to the corresponding pages for a leaf node, or the whole subtree under this node. While PDF/VT does not prescribe any particular kind of metadata, it is intended to be used with the JDF standard published by the CIP4 organization.

The CIP4 metadata format makes use of common PDF data types, and the metadata is written to the PDF output in a manner specified in the PDF/VT standard.

PDFlib users can use the POCA (PDF Object Creation API) interface to create arbitrary data structures consisting of PDF dictionaries, arrays, and other data types. The top-level dictionary for DPM metadata can be passed to the `dpm` option of `PDF_begin/end_dpart()`:

```

dpm          = p.poca_new("type=dict usage=dpm");
cip4_root    = p.poca_new("type=dict usage=dpm");
cip4_metadata = p.poca_new("type=dict usage=dpm");

p.poca_insert(dpm,          "type=dict  key=CIP4_Root value=" + cip4_root);
p.poca_insert(cip4_root,   "type=dict  key=CIP4_Metadata value=" + cip4_metadata);
p.poca_insert(cip4_metadata, "type=string key=CIP4_Conformance value=base");
p.poca_insert(cip4_metadata, "type=string key=CIP4_Creator value=starter_pdfvt1");
p.poca_insert(cip4_metadata, "type=string key=CIP4_JobID value={Kraxi Systems invoice}");

/* Create node in the DPart hierarchy and add DPM metadata */
p.begin_dpart("dpm=" + dpm);

p.poca_delete(dpm, "recursive=true");

```

Cookbook The PDFlib pCOS Cookbook includes a code sample for retrieving DPM from a PDF/VT document and creating the corresponding XML representation.

12.4.5 Scope Hints for recurring Graphical Content

Scope hints. The `scope` suboption of the `pdfvt` option should be supplied with suitable values to `PDF_load_image()`, `PDF_load_graphics()`, `PDF_open_pdi_page()`, and `PDF_begin_template_ext()`. In order to achieve this, the client application must know where and how often images, pages, and templates will be re-used in the current PDF document or (even better) across multiple documents. Although the `scope` option is not required, it is strongly recommended since it provides important optimization information for the RIP. If no such information is available, the `scope` option should be omitted or specified with the value `unknown` instead of supplying possibly wrong values.

For the scopes *singleuse*, *record* and *file* PDFlib checks whether the scope value is consistent with the actual use of the XObject, and emits a warning in the log file if the scope is too large, e.g. as follows:

```
XObject with handle 9 was assigned PDF/VT scope 'record', but was used only once
(use 'scope=singleuse' instead)
```

If you get such a warning you should check whether you can assign more suitable scope values to the corresponding image, graphic, imported PDF page, or template to avoid unnecessary caching in the RIP, thereby improving print performance.

While too few references to an XObject (relative to the *scope* option) trigger only a warning which indicates a possible optimization in client code, referencing an XObject too often is an error since this would violate the PDF/VT standard. More specifically, an exception will be thrown in the following situations:

- ▶ Option *scope=singleuse* was supplied and the XObject is used more than once in the document.
- ▶ Option *scope=record* was supplied and the XObject is used in more than one recipient record.

The *environment* suboption must be provided for *scope=stream* and *scope=global*. It should contain a suitable string which will be used to identify cached objects across documents. Depending on the workflow a customer or job reference may be used for this option. For example, a corporate logo which is heavily used across single print jobs may be identified this way.

Unique IDs. PDFlib automatically assigns unique IDs for all imported PDF pages, images and graphics (for graphics the *templateoptions* option must be provided). Equivalent images, pages, and graphics will be assigned identical ID values to enable efficient caching in the RIP. Unique IDs for templates should be provided by the user via the *xid* suboption of the *pdfvt* option. Identifiers for templates should be identical for template definitions which create equivalent PDF Form XObjects according to PDF/VT (i.e. templates with the same visual output). Templates which are not equivalent must have different identifiers or no identifier at all.

If the application needs the unique ID strings created by PDFlib, e.g. for use in DPM metadata, they can be retrieved with the *xid* option of *PDF_info_image()* for images and templates, *PDF_info_graphics()* for graphics, and *PDF_info_pdi_page()* for imported PDF pages.

12.4.6 Encapsulated XObjects

Encapsulating XObjects can greatly improve ripping performance. However, XObjects can only be encapsulated in fully opaque documents, or (in documents where transparency is used) if certain requirements are met.

Transparency in the document. Since PDFlib does not know in advance whether or not transparency will be used in the document the *usestransparency* option can be supplied to *PDF_begin_document()* as a hint. If the client uses this option to assert that no transparency will be used, all XObjects can be encapsulated. If transparency is used in the document the client must obey certain rules in order to have XObjects encapsulated. If *usestransparency=false* is supplied but the document contains transparent elements nevertheless PDFlib throws an exception.

Transparency is considered as being used in the generated document if one or more of the following conditions are true:

- ▶ `PDF_load_image()` or `PDF_fill_imageblock()` is called with any of the following options:
 - ▶ the image contains an alpha channel and the `ignoremask` option is `false`;
 - ▶ the `masked` option with a handle to an image with more than 1 bit per pixel.
- ▶ Graphics imported with `PDF_load_graphics()` contain any transparent element.
- ▶ `PDF_create_gstate()` is called with any of the following options:
 - ▶ the `softmask` option with an option list (i.e. not with the keyword `none`);
 - ▶ one of the options `opacityfill` or `opacitystroke` with a value different from 1;
 - ▶ the `blendmode` option with a value other than `None` and `Normal`.
- ▶ PDF pages imported with `PDF_open_pdi_page()` or `PDF_fill_pdfblock()` contain any transparent element. Transparency in imported PDF pages can be identified with the pCOS interface and the `usespagetransparency` pseudo object (see pCOS Path Reference for details).

Encapsulating XObjects. PDFlib encapsulates most XObjects according to the following rules:

- ▶ Image XObjects created with `PDF_load_image()` or `PDF_fill_imageblock()` are encapsulated if any of the following conditions is true:
 - ▶ the `renderingintent` option of `PDF_load_image()` has been supplied with a value different from `Auto`;
 - ▶ the `mask` option of `PDF_load_image()` has been supplied with the value `true`.
- ▶ Form XObjects created with `PDF_begin_template_ext()` are encapsulated if the document option `usestransparency=false` has been supplied, or the `transparencygroup` option has been supplied with the suboptions `colorspace` and `isolated=true`.
- ▶ Form XObjects created with `PDF_load_graphics()` or `PDF_load_image()` with the `templateoptions` option are always encapsulated. If the document option `usestransparency=true` has been supplied a transparency group with a suitable color space is attached to the Form XObject; otherwise no group is created.
- ▶ Form XObjects created with `PDF_open_pdi_page()` are encapsulated if the imported PDF page doesn't contain any layers or `PDF_open_pdi_document()` has been called with `uselayers=false`. If the document option `usestransparency=true` has been supplied a transparency group with a suitable color space is attached to the Form XObject; otherwise no group is created. If an imported page contains XObjects which are already encapsulated this property is retained unmodified when importing the page.

If an XObject cannot be encapsulated a warning is written to the log file.

Scope for encapsulated XObjects. It is recommended to supply the `scope` suboption of the `pdfvt` option for all encapsulated XObjects created with `PDF_load_image()`, `PDF_fill_imageblock()`, `PDF_load_graphics()` and `PDF_begin_template_ext()`. Since it is allowed to supply the `scope` suboption for unencapsulated XObjects as well, it can be supplied to all relevant API calls if the lifetime of an image, graphic, imported page or template is known, regardless of the encapsulation status of the generated XObject.

12.4.7 Importing PDF/X and PDF/VT Documents with PDI

Special rules apply when pages from an existing PDF document will be imported into a PDF/VT-conforming output document (see Section 8.3, »Importing PDF Pages with PDI«, page 180, for details on PDI). All imported documents must conform to either PDF/X-3/

4/4p or PDF/VT-1. If PDF/VT conformance is configured in PDFlib and the imported documents adhere to a compatible PDF/X level, the generated output is guaranteed to conform to the selected PDF/VT conformance level. Imported documents which do not adhere to one of the acceptable PDF/X levels are rejected.

DPart and DPM limitations with PDF/VT import. Note the following restriction when importing PDF/VT documents: Document part hierarchy (DPart) and document part metadata (DPM) are not imported. They can be queried with pCOS and reconstructed with *PDF_begin/end_dpart()* and POCA functions.



13 PPS and the PDFlib Block Plugin

The PDFlib Personalization Server (PPS) supports a template-driven PDF workflow for variable data processing. Using the Block concept, imported pages can be populated with variable amounts of single- or multi-line text, images, PDF pages or vector graphics. This can be used to easily implement applications which require customized PDF documents, for example:

- ▶ mail merge
- ▶ flexible direct mailings
- ▶ transactional and statement processing
- ▶ business card personalization

You can create and edit Blocks interactively with the PDFlib Block Plugin, convert existing PDF form fields to PDFlib Blocks with the form field conversion Plugin. Blocks can be filled with PPS. The results of Block filling with PPS can be previewed in Acrobat since the Block Plugin contains an integrated version of PPS.

Note Block processing requires the PDFlib Personalization Server (PPS). Although PPS is contained in all PDFlib packages, you must purchase a license key for PPS; a PDFlib or PDFlib+PDI license key is not sufficient. The PDFlib Block Plugin for Adobe Acrobat is required for creating Blocks in PDF templates interactively.

Cookbook Code samples regarding variable data and Blocks can be found in the blocks category of the PDFlib Cookbook.

13.1 Installing the PDFlib Block Plugin

The Block Plugin works with the following Acrobat versions (it doesn't work with Adobe Reader):

- ▶ Windows: Acrobat X/XI
- ▶ Windows: Acrobat Standard and Pro 2017/2020
- ▶ Windows: Acrobat Pro DC (32-bit and 64-bit)
- ▶ macOS: Acrobat Standard and Pro 2017/2020 (Intel)
- ▶ macOS: Acrobat DC (universal binary for Intel and M1)

Since Acrobat DC on Windows is available in 32-bit and 64-bit versions two different installers are available. It is important to use the appropriate installer which matches the installed Acrobat version.

Installing the PDFlib Block Plugin on Windows. To install the PDFlib Block Plugin and the PDF form field conversion Plugin in Acrobat, the plugin files must be placed in a subdirectory of the Acrobat plugin folder. This is done automatically by the plugin installer, but can also be done manually. The plugin files are called *Block.api* and *AcroFormConversion.api*.

The plugin folder for Acrobat 32-bit on 64-bit Windows typically looks as follows:

```
C:\Program Files (x86)\Adobe\Acrobat DC\Acrobat\plug_ins\PDFlib Block Plugin
```

The plugin folder for Acrobat 64-bit typically looks as follows:

```
C:\Program Files\Adobe\Acrobat DC\Acrobat\plug_ins\PDFlib Block Plugin
```

Installing the PDFlib Block Plugin for Acrobat on macOS. Proceed as follows to install the Plugin for all users:

- ▶ Double-click the disk image to mount it. A folder with the Plugin files will be visible.
- ▶ Copy the Plugin folder to the following path in the system's *Library* folder (create the *Plug-Ins* folder if it doesn't yet exist):

```
/Library/Application Support/Adobe/Acrobat/XXX/Plug-ins
```

Alternatively you can install the Plugin only for a single user as follows:

- ▶ Click the desktop to make sure you're in the Finder, hold down the *Option* key, and choose *Go, Library* to open the user's *Library* folder.
- ▶ Copy the Plugin folder to the following path in the user's *Library* folder (create the *Plug-Ins* folder if it doesn't yet exist):

```
/Users/<username>/Library/Application Support/Adobe/Acrobat/XXX/Plug-ins
```

Multi-lingual Interface. The PDFlib Block Plugin supports multiple languages in the user interface. Depending on the application language of Acrobat, the Block Plugin chooses its interface language automatically. Currently English, German and Japanese interfaces are available. If Acrobat runs in any other language mode, the Block Plugin uses the English interface.

Sandbox Protection for Acrobat DC on Windows. Acrobat DC 2020 introduced a new security model called *Sandbox Protections* which can be activated via *Preferences, Security (Enhanced), Protected Mode* and *Protected View*. If it is enabled various operations are restricted and a yellow bar with a security message appears at the top of the document window. More information about Sandbox Protections can be found at:

helpx.adobe.com/acrobat/using/whats-new/2020-august.html

www.adobe.com/devnet-docs/acrobatetk/tools/AppSec/sandboxprotections.html

If Sandboxing is enabled it affects the Preview feature of the PDFlib Block Plugin. Protected View by default grants access to Acrobat's *AppData* directory, the temp directory and several other directories, but not to arbitrary user directories. The Block Plugin can only read from and write to directories which are included in the default directory list of Protected View or which have been configured (whitelisted) in a policy file at the following location (for 32-bit and 64-bit versions of Acrobat):

```
C:\Program Files (x86)\Adobe\Acrobat DC\Acrobat\PDFlibBlockCustomPolicies.txt
```

```
C:\Program Files\Adobe\Acrobat DC\Acrobat\PDFlibBlockCustomPolicies.txt
```

By default the policy file grants access to the following directories, but more directory names can be added by the Administrator:

```
; Protected Path Section  
FILES_ALLOW_ANY = C:\Users\<username>\*.*  
FILES_ALLOW_ANY = C:\Users\Public\*.*
```

If Protected Mode or Protected View is enabled and directories are used which are not whitelisted, some features of the Block Plugin including Preview and Block import/export may fail.

Troubleshooting. If the PDFlib Block Plugin doesn't seem to work check the following:

- ▶ Make sure that in *Edit, Preferences, [General...], General* the box *Use only certified plug-ins* is unchecked. The plugins are not loaded if Acrobat runs in Certified Mode.
- ▶ Some PDF forms created with Adobe Designer or Adobe Experience Manager may prevent the Block Plugin as well as other Acrobat plugins from working properly since they interfere with Acrobat's internal security model. For this reason we suggest to avoid Designer's static PDF forms, and only use dynamic PDF forms as input for the Block Plugin.

13.2 Overview of the Block Concept

13.2.1 Separation of Document Design and Program Code

PDFlib Blocks make it easy to place variable text, images, PDF pages or vector graphics on imported pages. In contrast to simple PDF pages, pages with Blocks intrinsically carry information about the required processing which will be performed later on the server side. The Block concept separates the following tasks:

- ▶ The designer creates the page layout and specifies the location of variable page elements along with relevant properties such as font size, color, or image scaling. After creating the layout as a PDF document, the designer uses the PDFlib Block Plugin for Acrobat to specify variable data Blocks and their associated properties.
- ▶ The programmer writes code to connect the information contained in PDFlib Blocks on imported PDF pages with dynamic information, e.g., database fields. The programmer doesn't need to know any details about a Block (whether it contains a name or a ZIP code, the exact location on the page, its formatting, etc.) and is therefore independent from any layout changes. PPS will take care of all Block-related details based on the Block properties found in the file.

In other words, the code written by the programmer is »data-blind« – it is generic and does not depend on the particulars of any Block. For example, the designer can move the Block with name of the addressee in a mailing to a different location on the page, or change the font size. The generic Block handling code doesn't need to be changed, and will generate correct output once the designer changed the Block properties with the Acrobat plugin to use the first name instead of the last name.

As an intermediate step Block filling can be previewed in Acrobat to accelerate the development and test cycles. Block previews are based on default data (e.g. a string or an image file name) which is specified in the Block definitions.

13.2.2 Block Properties

The behavior of Blocks can be controlled with Block properties. Properties are assigned to a Block with the Block Plugin.

Predefined Block properties. Blocks are defined as rectangles on the page which are assigned a name, a type, and an open set of properties which will later be processed by PPS. The name is an arbitrary string which identifies the Block, such as *firstname*, *lastname*, or *zipcode*. PPS supports different kinds of Blocks:

- ▶ *Textline Blocks* hold a single line of textual data which will be processed with the Textline output method in PPS.

- ▶ *Textflow Blocks* hold one or more lines of textual data. Multi-line text will be formatted with the Textflow formatter in PPS. Textflow Blocks can be linked so that one Block holds the overflow text of the previous Block (see »Linking Textflow Blocks«, page 355).
- ▶ *Image Blocks* hold a raster image. This is similar to placing a TIFF or JPEG file in a DTP application.
- ▶ *PDF Blocks* hold arbitrary PDF contents imported from a page in another PDF document. This is similar to placing a PDF page in a DTP application.
- ▶ *Graphics Blocks* hold vector graphics. This is similar to placing an SVG file in a layout application.

Blocks can carry a number of properties depending on their type. Properties can be created and modified with the Block Plugin (see Section 13.3.2, »Editing Block Properties«, page 341). A full list of Block properties can be found in Section 13.7, »Block Properties«, page 357. For example, a text Block can specify the font and size of the text, an image or PDF Block can specify the scaling factor or rotation PPS offers dedicated functions for processing the Block types, e.g. `PDF_fill_textblock()`. These functions search a placed PDF page for a Block by its name, analyze its properties, and place client-supplied data (single- or multi-line text, raster image, PDF page, or vector graphics) on the new page according to the specified Block properties. The programmer can override Block properties by specifying the corresponding options in the Block filling functions.

Properties for default contents. Special Block properties can be defined which hold the default contents of a Block, i.e. the text, image, PDF, or graphics contents that will be placed in the Block if no variable data is supplied to the Block filling functions, or in situations where the Block contents are currently constant, but may change in the next print run.

Default properties are also used by the Preview feature of the Block Plugin (see Section 13.5, »Previewing Blocks in Acrobat«, page 349).

Custom Block properties. Predefined Block properties make it possible to quickly implement variable data processing applications, but they are restricted to the set of properties which are internally known to PPS and can automatically be processed. In order to provide more flexibility, the designer can also assign custom properties to a Block. These can be used to extend the Block concept in order to match the requirements of more advanced variable data processing applications.

There are no rules for custom properties since PPS will not process custom properties in any way, except making them available to the client. The client code can retrieve the value of custom properties and process it as appropriate. Based on a custom property of a Block the application may make layout-related or data-gathering decisions. For example, a custom property for a scientific application could specify the number of digits for numerical output, or a database field name may be defined as a custom Block property for retrieving the data corresponding to this Block.

13.2.3 Why not use PDF Form Fields?

Experienced Acrobat users may ask why we implemented a new Block concept instead of relying on the existing form field mechanism available in PDF. The primary distinction is that PDF form fields are optimized for interactive filling, while PDFlib Blocks are targeted at automated filling. Applications which need both interactive and automated

filling can combine PDF forms and PDFlib Blocks with the form field conversion plugin (see Section 13.4, »Converting PDF Form Fields to PDFlib Blocks«, page 346).

Although there are many parallels between both concepts, PDFlib Blocks offer several advantages over PDF form fields as detailed in Table 13.1.

Table 13.1 Comparison of PDF form fields and PDFlib Blocks

<i>feature</i>	<i>PDF form fields</i>	<i>PDFlib Blocks</i>
<i>design objective</i>	<i>for interactive use</i>	<i>for automated filling</i>
<i>typographic features (beyond choice of font and font size)</i>	–	<i> Kerning, word and character spacing, underline/overline/strikeout</i>
<i>OpenType layout features</i>	–	<i>dozens of OpenType layout features, e.g. ligatures, swash characters, oldstyle figures</i>
<i>complex script support</i>	<i>limited</i>	<i>shaping and bidirectional formatting, e.g. for Arabic and Devanagari</i>
<i>font control</i>	<i>font embedding</i>	<i>font embedding and subsetting, encoding</i>
<i>text formatting controls</i>	<i>left-, center-, right-aligned</i>	<i>left-, center-, right-aligned, justified; various formatting algorithms and controls; inline options can be used to control the appearance of text</i>
<i>change font or other text attributes within text</i>	–	<i>yes</i>
<i>merged result is integral part of PDF page description</i>	–	<i>yes</i>
<i>users can edit merged field contents</i>	<i>yes</i>	<i>no</i>
<i>extensible set of properties</i>	–	<i>yes (custom Block properties)</i>
<i>use image files for filling</i>	–	<i>BMP, PNG, JPEG, JBIG2, JPEG 2000, TIFF</i>
<i>use vector graphics for filling</i>	–	<i>SVG</i>
<i>color support</i>	<i>RGB</i>	<i>grayscale, RGB, CMYK, Lab, spot color (HKS and Pantone spot colors integrated in the Block Plugin), DeviceN</i>
<i>PDF standards</i>	–	<i>PDF/A, PDF/X, PDF/VT, PDF/UA</i>
<i>graphics and text properties can be overridden upon filling</i>	–	<i>yes</i>
<i>transparent contents</i>	–	<i>yes</i>
<i>Text Blocks can be linked (text flows from one Block to the next)</i>	–	<i>yes</i>

13.3 Editing Blocks with the Block Plugin

13.3.1 Creating Blocks

Activating the Block tool. The Block Plugin for creating PDFlib Blocks is similar to the form tool in Acrobat. All Blocks on the page will be visible when the Block tool is active. When another Acrobat tool is selected the Blocks will be hidden, but they are still present. You can activate the Block tool in the following ways:

- ▶ By clicking the Block icon  which you can locate as follows:
 - Acrobat DC: click *Tools, Advanced Editing*.
- ▶ Via the menu item *PDFlib Blocks, PDFlib Block Tool*.

Creating and modifying Blocks. When the Block tool is active you can drag the cross-hair pointer to create a Block at the desired position on the page and with the desired size. Blocks are always rectangular with edges parallel to the page edges (use the *rotate* property for Block contents which are not parallel to the page edges). After dragging a Block rectangle the Block properties dialog appears where you can edit the properties of the Block (see Section 13.3.2, »Editing Block Properties«, page 341). The Block tool automatically creates a Block name which can be changed in the properties dialog. Block names must be unique on a page, but can be repeated on another page.

You can change the Block type in the top area to one of *Textline, Textflow, Image, PDF, or Graphics*. Different colors are used for representing the Block types (see figure 13.1). The *Block Properties* dialog hierarchically organizes the properties in groups and sub-groups depending on the Block type.

Note After you added Blocks or made changes to existing Blocks in a PDF, use Acrobat's »Save as...« command (as opposed to »Save«) to achieve smaller file sizes.

Selecting Blocks. Several Block operations, such as copying, moving, deleting, or editing Properties, work with one or more selected Blocks. You can select Blocks with the Block tool as follows:

- ▶ To select a single Block simply click on it.
- ▶ To select multiple Blocks hold down the Shift key while clicking on the second and subsequent Block.
- ▶ Press Ctrl-A (on Windows) or Cmd-A (on macOS) or *Edit, Select All* to select all Blocks on a page.

The context menu. When one or more Blocks are selected you can open the context menu to quickly access Block-related functions (which are also available in the *PDFlib Blocks* menu). To open the context menu, click on the selected Block(s) with the right mouse button on Windows, or Ctrl-click the Block(s) on macOS. For example, to delete a Block, select it with the Block tool and press the *Delete* key, or use *Edit, Delete* in the context menu.

If you right-click (or Ctrl-click on macOS) an area on the page where no Block is located the context menu contains entries for creating a Block Preview and for configuring the Preview feature.

Block size and position. Using the Block tool you can move one or more selected Blocks to a different position. Hold down the Shift key while dragging a Block to restrain the positioning to horizontal and vertical movements. This may be useful for exactly aligning Blocks. When the pointer is located near a Block corner, the pointer icon changes to a double arrow and you can resize the Block. The Block width and height are displayed after the Block as long as the cursor is located inside the Block area.

To adjust the position or size of multiple Blocks, select two or more Blocks and use the *Align, Center, Distribute, or Size* commands from the *PDFlib Blocks* menu or the context menu. The position of one or more Blocks can also be changed in small increments (one point) by using the arrow keys. If you press the *Shift* key while moving Blocks with the

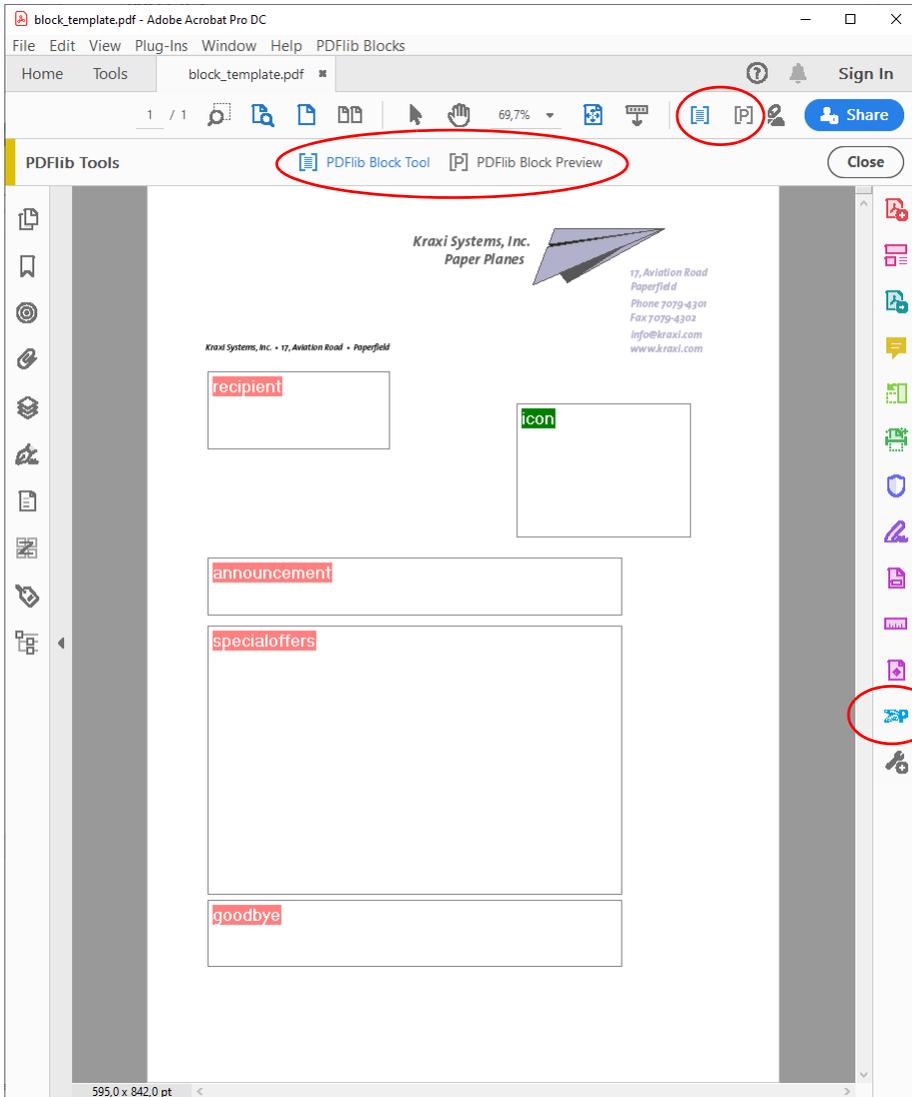


Fig. 13.1
Visualization
of Blocks

arrow keys the increment is 10 points. Since the increment is converted to screen pixels the actual change of the Block rectangle coordinates may be slightly different from 1 or 10 points, though.

Alternatively, you can enter numerical Block coordinates in the properties dialog. The origin of the coordinate system is in the upper left corner of the page. The coordinates will be displayed in the unit which is currently selected in Acrobat:

- ▶ To change the display units in Acrobat DC proceed as follows: go to *Edit, Preferences, [General...], Units & Guides, Page & Ruler Units* and choose one of Points, Inches, Millimeters, Picas, Centimeters.
- ▶ To display cursor coordinates use *View, Show/Hide, Cursor Coordinates*.

Note that the selected unit affects only the *Rect* property, but not any other numerical properties (e.g. *fontsize*).

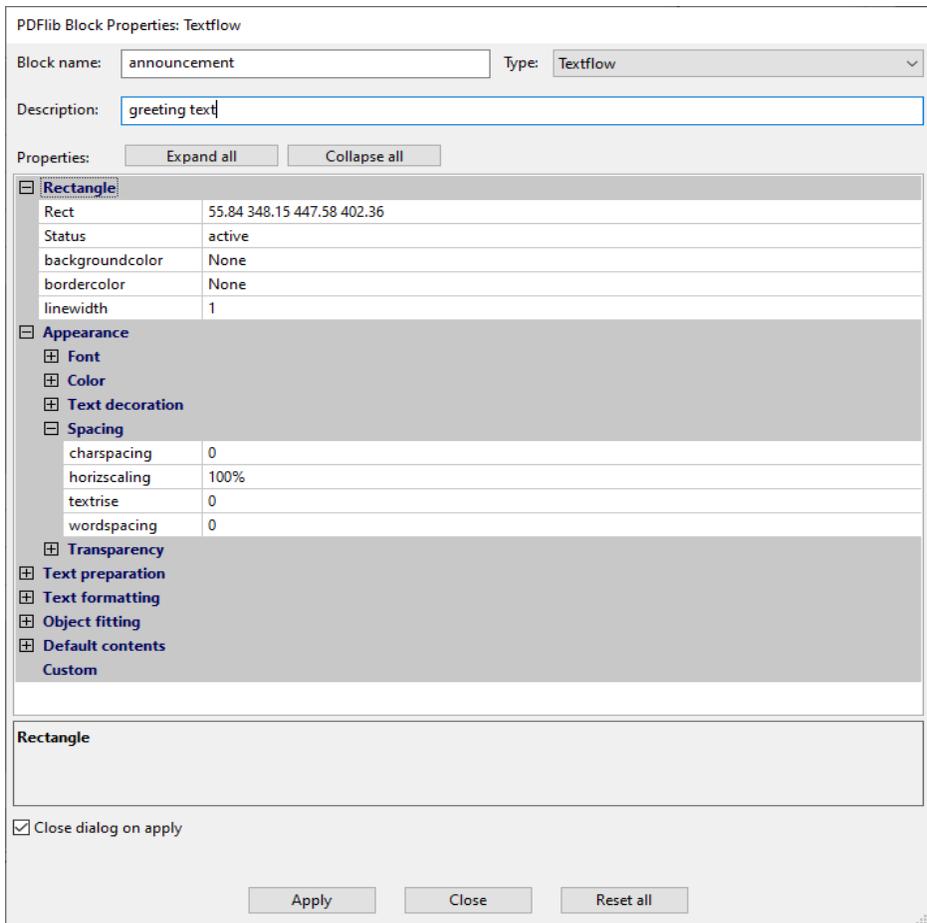


Fig. 13.2
The Block properties dialog

Using a grid to position Blocks. You can take advantage of Acrobat's grid feature for precisely positioning and resizing Blocks:

- ▶ Display the grid: *View, Show/Hide, Rulers & Grids, Grid*;
- ▶ Enable grid snapping: *View, Show/Hide, Rulers & Grids, Snap to Grid*;
- ▶ Change the grid: go to *Edit, Preferences, [General...], Units & Guides*. Here you can change the spacing and position of the grid as well as the color of the grid lines.

If *Snap to Grid* is enabled the size and position of Blocks will be aligned with the configured grid. *Snap to Grid* affects newly generated Blocks as well as existing Blocks which are moved or resized with the Block tool.

Creating Blocks by selecting an image or graphic. As an alternative to manually dragging Block rectangles you can use existing page contents to define the Block size. First, make sure that the menu item *PDFlib Blocks, Click Object to define Block* is enabled. Now you can use the Block tool to click on an image on the page in order to create a Block with the same size and location as the image. You can also click on other graphical ob-

jects, and the Block tool will try to select the surrounding graphic (e.g., a logo). The *Click Object* feature is intended as an aid for defining Blocks. If you want to reposition or resize the Block you can do so afterwards without any restriction. The Block will not be locked to the image or graphics object which was used as a positioning aid.

The *Click Object* feature tries to recognize which vector graphics and images form a logical element on the page. When some page content is clicked, its bounding box (the surrounding rectangle) will be selected unless the object is white or very large. In the next step other objects which are partially contained in the detected rectangle will be added to the selected area, and so on. The final area will be used as the basis for the generated Block rectangle. The end result is that the *Click Object* feature will try to select complete graphics, not only individual lines.

Automatically detect font properties. The Block Plugin can analyze the underlying font which is present at the location where a Textline or Textflow Block is positioned, and can automatically fill in the corresponding properties of the Block:

fontname, fontsize, fillcolor, charspacing, horzscaling, wordspacing, textrendering, textrise

Since automatic detection of font properties can result in undesired behavior if the background shall be ignored, it can be activated or deactivated using *PDFlib Blocks, Detect underlying font and color*. By default this feature is turned off.

Locking Blocks. Blocks can be locked to protect them against accidentally moving, resizing, or deleting. With the Block tool active, select the Block and choose *Lock* from its context menu. While a Block is locked you cannot move, resize, or delete it, nor edit its properties.

13.3.2 Editing Block Properties

When you create a new Block, double-click an existing one, or choose *Properties* from a Block's context menu, the properties dialog will appear where you can edit all settings related to the selected Block (see figure 13.2). As detailed in Section 13.7, »Block Properties«, page 357, there are several groups of properties available, subject to the Block type. The *Apply* button will only be enabled if you changed one or more properties in the dialog. The *Apply* button will be inactive for locked Blocks.

Note Some properties may be inactive depending on the Block type and certain property settings. For example, the property subgroup *Ruler tabs* for `hortabmethod=ruler` where you can edit *tabulator settings* is enabled only if the *hortabmethod* property in the group *Text formatting, Tabs* is set to *ruler*.

Note If you enter text for a Block property on macOS you may experience character replacements, e.g. straight quotes are replaced by smart quotes. This substitution is done by the operation system and can be disabled via »System Preferences«, »Keyboard«, »Text«, »Use smart quotes and dashes«.

To change a property's value enter the desired number or string in the property's input area (e.g. *linewidth*), choose a value from a drop-down list (e.g. *fitmethod*, *orientate*), or select a font, color value or file name by clicking the »...« button at the right-hand side of the dialog (e.g. *backgroundcolor*, *defaultimage*). For the *fontname* property you can either choose from the list of fonts installed on the system or type a custom font name. Re-

ardless of the method for entering a font name, the font must be available on the system where the Blocks will be filled with PPS.

Modified properties will in be displayed in bold face in the Block Properties dialog. If any of the properties in a Block has been modified, the suffix (*) will be appended to the displayed Block name. When you are done editing properties click the *Apply* button to update the Block. The properties just defined will be stored in the PDF file as part of the Block definition.

Stacked Blocks. Overlapping Blocks can be difficult to select since clicking an area will always select the topmost Block. In this situation the *Select Block* entry in the context menu can be used to select one of the Blocks by name. As soon as a Block has been selected this way, the next action within its area will not affect other Blocks, but only the selected one. For example, press *Enter* to edit the selected Block's properties. This way Block properties can easily be edited even for Blocks which are partially or completely covered by other Blocks.

Using and restoring repeated values of Block properties. In order to save some amount of typing and clicking, the Block tool remembers the property values which have been entered into the previous Block's properties dialog. These values will be reused when you create a new Block. Of course you can override these values with different ones at any time.

Pressing the *Reset all* button in the properties dialog resets most Block properties to their respective default values. The following items remain unmodified:

- ▶ the *Name*, *Type*, *Rect*, and *Description* properties;
- ▶ all custom properties.

Note Do not confuse the default values of predefined Block properties with the defaulttext, defaultimage, defaultpdf, and defaultgraphics properties which hold placeholder data for generating previews (see »Default Block contents«, page 349).

Editing multiple Blocks at once. Editing the properties of multiple Blocks at once is a big time saver. You can select multiple Blocks as follows:

- ▶ Activate the Block tool via the menu item *PDFlib Blocks, PDFlib Block Tool*.
- ▶ Click on the first Block to select it. The first selected Block is the master Block. Shift-click other Blocks to add them to the set of selected Blocks. Alternatively, click *Edit, Select All* to select all Blocks on the current page.
- ▶ Double-click on any of the Blocks to open the Block Properties dialog. The Block where you double-click will be the new master Block.
- ▶ Alternatively, you can click on a single Block to designate it as master Block, and then press the *Enter* key to open the Block Properties dialog.

The Properties dialog displays only the subset of properties which apply to all selected Blocks. The dialog is populated with property values taken from the master Block. Now you can apply properties to all selected Blocks as follows:

- ▶ If the checkbox *Apply all properties of the master Block* is unchecked: upon clicking *Apply* only the properties changed manually in the dialog (highlighted in black) are copied to all selected Blocks.
- ▶ If the checkbox *Apply all properties of the master Block* is checked: upon pressing *Apply* all current properties of the master Block as well as all properties changed manually

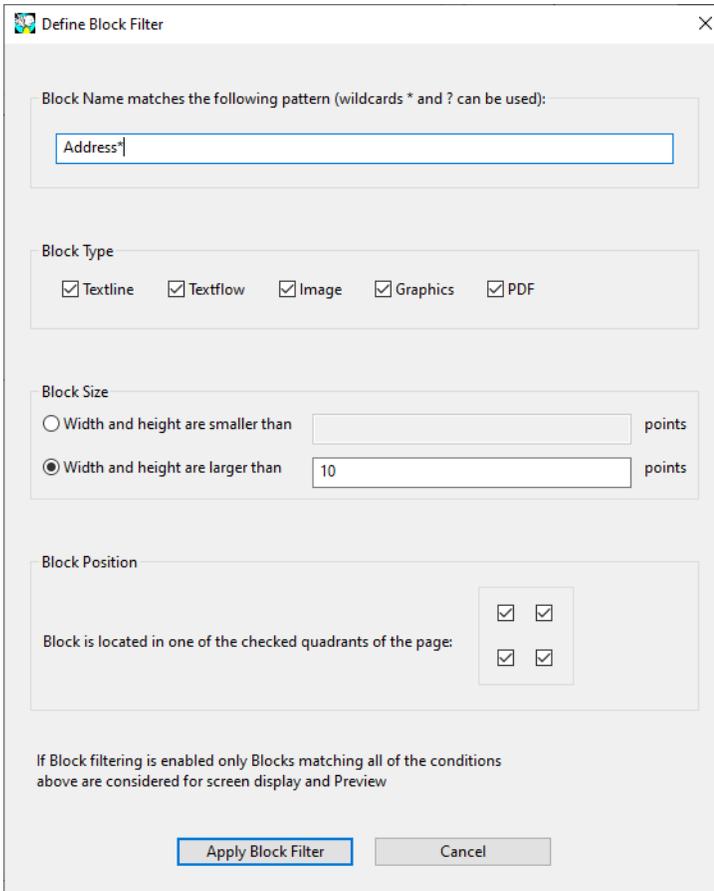


Fig. 13.3
Block Filter definition

in the dialog are copied to all selected Blocks. This can be used to copy all properties from a particular Block to one or more other Blocks.

The following predefined properties as well as custom properties can not be shared, i.e. they can not be edited for multiple Blocks at once:

Name, Description, Subtype, Type, Rect, Status

13.3.3 Block Filtering

If a page contains a large number of Blocks editing the Blocks may become inconvenient because of many Block rectangles and because screen display performance suffers. In this situation you can use Block filtering which restricts screen display to a subset of Blocks based on one or more of the following criteria which can be specified in the *Define Block Filter* dialog (see figure 13.3):

- ▶ Block name: you can select Blocks where the name matches the specified pattern; the pattern may include the wildcards * (any character sequence) and ? (any single character);
- ▶ Block type;
- ▶ minimum and maximum Block width/height;

- ▶ location of the Block in one of the page quadrants, i.e. upper left quarter of the page, lower left etc.

Once a filter has been defined it can be enabled and disabled via the *Apply Block Filter* item in the context menu. Block filtering affects screen display and Preview (see Section 13.5, »Previewing Blocks in Acrobat«, page 349). Block filter definition and activation state are maintained separately for each document. Active Block filtering is visualized with a small »F« icon in the top left corner of the page as long as the Block tool is active.

The last used filter definition is saved in the configuration file when Acrobat is closed, and restored upon the next restart. To avoid confusion Block filtering is disabled when a document is opened.

13.3.4 Copying Blocks between Pages and Documents

The Block Plugin offers several methods for moving and copying Blocks within the current page, the current document, or between documents:

- ▶ move or copy Blocks by dragging them with the mouse, or pasting Blocks to another page or open document
- ▶ duplicate Blocks on one or more pages of the same document using standard copy/paste operations
- ▶ export Blocks to a new file (with empty pages) or to an existing document (apply the Blocks to existing pages)
- ▶ import Blocks from another document

In order to update the page contents while maintaining Block definitions you can replace the underlying page(s) while keeping the Blocks. Use *Tools, Organize Pages, Replace* for this purpose.

Copying Blocks. You can copy Blocks by selecting one or more Blocks and dragging them to a new location while pressing the Ctrl key (on Windows) or Alt key (on macOS). The accelerator keys are configurable (see Section 13.3.5, »Customizing the Block Plugin User Interface with XML«, page 345). The mouse cursor changes while this key is pressed. A copied Block has the same properties as the original Block with the exception of its name and position which are automatically adjusted in the new Block.

You can also use copy/paste to copy Blocks to another location on the same page, to another page in the same document, or to another document which is currently open in Acrobat:

- ▶ Activate the Block tool and select the Blocks you want to copy.
- ▶ Use Ctrl-C (on Windows) or Cmd-C (on macOS) or *Edit, Copy* to copy the selected Blocks to the clipboard.
- ▶ Navigate to the target page (if necessary).
- ▶ Make sure that the Block Tool is active, and use Ctrl-V (on Windows) or Cmd-V (on macOS) or *Edit, Paste* to paste the Blocks from the clipboard to the current page and document.

Duplicating Blocks on other pages. You can create duplicates of one or more Blocks on an arbitrary number of pages in the current document simultaneously:

- ▶ Activate the Block tool and select the Blocks you want to duplicate.
- ▶ Choose *Import and Export, Duplicate...* from the *PDFlib Blocks* menu or the context menu.

- ▶ Choose which Blocks to duplicate (*Selected Blocks* or *All Blocks on this Page*) and the range of target pages to which you want to duplicate the selected Blocks.

Exporting and importing Blocks. Using the export/import feature for Blocks it is possible to share the Block definitions on a single page or all Blocks in a document among multiple PDF files. This is useful for updating the page contents while maintaining existing Block definitions. To export Block definitions to a separate file proceed as follows:

- ▶ Activate the Block tool and select the Blocks you want to export.
- ▶ Choose *Import and Export, Export...* from the *PDFlib Blocks* menu or the context menu. Enter the page range and a file name of the new PDF with the Block definitions.

You can import Block definitions via *PDFlib Blocks, Import and Export, Import...* Upon importing Blocks you can choose whether to apply the imported Blocks to all pages in the document or only to a page range. If more than one page is selected the Block definitions will be copied unmodified to the pages. If there are more pages in the target range than in the imported Block definition file you can use the *Repeat Template* checkbox. If it is enabled the sequence of Blocks in the imported file will be repeated in the current document until the end of the document is reached.

Copying Blocks to another document upon export. When exporting Blocks you can immediately apply them to the pages in another document, thereby propagating the Blocks from one document to another. In order to do so choose an existing document to export the Blocks to. If you activate the checkbox *Delete existing Blocks* all Blocks which may be present in the target document will be deleted before copying the new Blocks into the document.

13.3.5 Customizing the Block Plugin User Interface with XML

Some aspects of the Block Plugin user interface can be controlled via an XML configuration file. A sample configuration file *factory settings.xml* is included in the distribution. If the configuration has been modified the new settings are stored in *user settings.xml*. The modified configuration will be loaded every time Acrobat is started and written when Acrobat is closed. The configuration file is stored in a location similar to the following (the names of system directories may be localized):

Windows: C:\Users\<user>\AppData\Local\Adobe\Acrobat\DC\PDFlib\Block Plugin 6
macOS: /Users/<user>/Library/Application Support/Adobe/Acrobat/DC/PDFlib/Block Plugin 6

The following XML elements can be used to modify the configuration manually:

- ▶ The element */Block_Plugin/MainDialog/CloseOnApply* controls the initial status of the *Close dialog on apply* checkbox in the *Block Properties* dialog. This checkbox determines whether the dialog will be kept open after creating a Block or modifying Block properties.
- ▶ The element */Block_Plugin/MainDialog/ApplyAllProps* controls the initial status of the *Apply all properties of the master Block* checkbox in the *Block Properties* dialog. This checkbox determines whether all properties of the master Block are copied to multiple selected Blocks or only those properties which have been modified in the dialog.
- ▶ The element */Block_Plugin/FontDialog/ShowBaseFonts* controls whether the base 14 fonts will be displayed in the font list of the Block Properties dialog (property group *Appearance*, property *fontname*) in addition to the fonts installed on the system.

- ▶ The element `/Block_Plugin/Command/ControlByClick` controls the initial status of the menu item *PDFlib Blocks, Click object to define Block*.
- ▶ The element `/Block_Plugin/Command/DetectFonts` controls the initial status of the menu item *PDFlib Blocks, Detect underlying font and color*.
- ▶ The element `/Block_Plugin/Command/KeyAccelerator` with the possible values *control* (which designates the Ctrl key on Windows and the Command key on macOS), *shift* for the Shift key, *control+shift* or *none* specifies the accelerator key for the following keyboard shortcuts:

A (select all), C (copy), I (Block Properties dialog), V (paste), X (cut)

The change will be effective upon the next start of Acrobat since keyboard shortcuts cannot be changed at runtime. If this entry is absent, no accelerators are available. The default is *control*.

- ▶ The element `/Block_Plugin/Configuration/Preferences/PreviewStatusMessage` controls whether a status message dialog (e.g. »10 Block(s) processed: ...«) is shown after each Preview operation.
- ▶ The element `/Block_Plugin/Service/MoveAccelerator` accepts the values *shift* and *control* which specify the modifier key for fast Block moving with the array keys. The element `/Block_Plugin/Service/MoveAcceleratorSize` specifies the increment when the modifier key is pressed (default: 10 points).

13.4 Converting PDF Form Fields to PDFlib Blocks

As an alternative to creating PDFlib Blocks manually, you can automatically convert PDF form fields to Blocks. This is especially convenient if you have complex PDF forms which you want to fill automatically with PPS or need to convert a large number of existing PDF forms for automated filling. In order to convert all form fields on a page to PDFlib Blocks choose *PDFlib Blocks, Convert Form Fields, Current Page*. To convert all form fields in a document choose *All Pages* instead. Finally, you can convert only selected form fields (choose Acrobat's *Select Object Tool* via *Tools, Rich Media*) to select one or more form fields) with *Selected Form Fields*.

Form field conversion details. Automatic form field conversion will convert form fields of the types selected in the *PDFlib Blocks, Convert Form Fields, Conversion Options...* dialog to Blocks of type *Textline* or *Textflow*. By default all form field types will be converted. Attributes of the converted fields will be transformed to the corresponding Block properties according to Table 13.3.

Multiple form fields with the same name. Multiple form fields on the same page are allowed to have the same name, while Block names must be unique on a page. When converting form fields to Blocks a numerical suffix will therefore be added to the name of generated Blocks in order to create unique Block names (see also »Associating form fields with corresponding Blocks«, page 347).

Note that due to a problem in Acrobat the field attributes of form fields with the same names are not reported correctly. If multiple fields have the same name, but different attributes these differences will not be reflected in the generated Blocks. The Conversion process will issue a warning in this case and provide the names of affected form fields. In this case you should carefully check the properties of the generated Blocks.

Associating form fields with corresponding Blocks. Since the form field names will be modified when converting multiple fields with the same name (e.g. radio buttons) it is difficult to reliably identify the Block which corresponds to a particular form field. This is especially important when using an FDF or XPDF file as the source for filling Blocks such that the final result resembles the filled form.

In order to solve this problem the AcroFormConversion plugin records details about the original form field as custom properties when creating the corresponding Block. Table 13.2 lists the custom properties which can be used to reliably identify the Blocks; all properties have type *string*.

Table 13.2 Custom properties for identifying the original form field corresponding to the Block

<i>custom property</i>	<i>meaning</i>
<i>PDFlib:field:name</i>	Fully qualified name of the form field
<i>PDFlib:field:pagenumber</i>	Page number (as a string) in the original document where the form field was located
<i>PDFlib:field:type</i>	Type of the form field; one of pushbutton, checkbox, radiobutton, listbox, combobox, textfield, signature
<i>PDFlib:field:value</i>	(Only for type=checkbox) Export value of the form field

Binding Blocks to the corresponding form fields. In order to keep PDF form fields and the generated PDFlib Blocks synchronized, the generated Blocks can be bound to the corresponding form fields. This means that the Plugin internally maintains the relationship of form fields and Blocks. When the conversion process is activated again, bound Blocks are updated to reflect the attributes of the corresponding PDF form fields. Bound Blocks are useful to avoid duplicate work: when a form is updated for interactive use, the corresponding Blocks can automatically be updated, too.

If you do not want to keep the converted form fields after Blocks have been generated you can choose the option *Delete converted Form Fields* in the *PDFlib Blocks, Convert Form Fields, Conversion Options...* dialog. This option will permanently remove the form fields after the conversion process. Any actions (e.g., JavaScript) associated with the affected fields will also be removed from the document.

Batch conversion. If you have many PDF documents with form fields that you want to convert to PDFlib Blocks you can automatically process an arbitrary number of documents using the batch conversion feature. The batch processing dialog is available via *PDFlib Blocks, Convert Form Fields, Batch conversion...*:

- ▶ The input files can be selected individually; alternatively the full contents of a folder can be processed.
- ▶ The output files can be written to the same folder where the input files are, or to a different folder. The output files can receive a prefix to their name in order to distinguish them from the input files.
- ▶ When processing a large number of documents it is recommended to specify a log file. After the conversion it will contain a full list of processed files as well as details regarding the result of each conversion along with possible error messages.

During the conversion process the converted PDF documents are visible in Acrobat, but you cannot use Acrobat's menu functions or tools until the conversion is finished.

Table 13.3 Conversion of PDF form fields to PDFlib Blocks

PDF form field attribute...	...will be converted to the PDFlib Block property
all fields	
Position	Rect
Name	Name
Tooltip	Description
Appearance, Text, Font	fontname
Appearance, Text, Font Size	fontsize; auto font size will be converted to a fixed font size of 2/3 of the Block height, and fitmethod will be set to auto. For multi-line fields/Blocks this combination will automatically result in a suitable font size which may be smaller than the initial value of 2/3 of the Block height.
Appearance, Text, Text Color	strokecolor and fillcolor
Appearance, Border, Border Color	bordercolor
Appearance, Border, Fill Color	backgroundcolor
Appearance, Border, Line Thickness	linewidth: Thin=1, Medium=2, Thick=3
General, Common Properties, Form Field	Status: Visible=active Hidden=ignore Visible but doesn't print=ignore Hidden but printable=active
General, Common Properties, Orientation	orientate: 0=north, 90=west, 180=south, 270=east
text fields	
Options, Default Value	defaulttext
Options, Alignment	position: Left={left center} Center={center center} Right={right center}
Options, Multi-line	checked creates Textflow Block unchecked creates a Textline Block
radio buttons and check boxes	
If »Check box/Button is checked by default« is selected: Options, Check Box Style or Options, Button Style	defaulttext: Check=4 Circle=l Cross=8 Diamond=u Square=n Star=H (these characters represent the respective symbols in the ZapfDingbats font)
list boxes and combo boxes	
Options, Selected (default) item	defaulttext
buttons	
Options, Icon and Label, Label	defaulttext

13.5 Previewing Blocks in Acrobat

Note You can try the Preview feature with the `block_template.pdf` document in the PDFlib distribution. The required resources (e.g. font and image) are also included in the PDFlib distribution.

PDFlib Blocks will be processed by PPS where the Block filling process can be customized regarding the data sources (e.g. text from a database, image files on disk) as well as visual and interactive aspects of the generated documents. This process is detailed in Section 13.6, »Filling Blocks with PPS«, page 353.

However, the Block Plugin contains an integrated version of PPS which can be used to generate Preview versions of the filled Blocks interactively in Acrobat without any programming. Although this Preview feature cannot offer the same flexibility as custom programming, it provides a quick overview of Block filling results. The Block Preview can be used for improving the position and size of Blocks as well as for checking the Block properties (e.g. font name and size). You can change the Blocks and create a new Preview until you are satisfied with the results shown in the Preview. Previews can be generated for the current page or the whole document.

The Preview will always be shown in a new PDF document. The original document (which contains the Blocks) will not be modified by generating a Preview. You can save or discard the generated Preview documents according to your requirements.

Default Block contents. Since the server-side data sources (e.g. a database) for the text, image, vector graphics or PDF contents of a Block are not available in the Plugin, the Preview feature always uses the Block's default contents, i.e. the values specified in the `defaulttext`, `defaultimage`, `defaultpdf`, or `defaultgraphics` properties. Usually, a sample data set will be used as default data which is representative for the real Block contents used with PPS. Blocks without any default contents are ignored when generating the Preview, as well as Blocks with `Status=ignoredefault`.

The default properties are empty for new Blocks. Before using the Preview feature you must fill the `defaulttext`, `defaultimage`, `defaultpdf`, or `defaultgraphics` properties (depending on the Block type) in the *Default contents* property group, or supply suitable values for the options of the same name in the *Advanced PPS options...* dialog.

Generating Block Previews. You can create Block Previews with one of the following methods:

- ▶ By clicking the PDFlib Block Preview icon  which you can locate as follows in Acrobat DC: click *Tools, Advanced Editing*.
- ▶ Via the menu item *PDFlib Blocks, Preview, Generate Preview*.
- ▶ If the Block tool is active you can right-click outside of any Block to bring up a context menu with the entries *Generate Preview* and *Preview Configuration*.

The Previews will be created based on the PDF file on disk. Any changes that you may have applied in Acrobat will only be reflected in the Preview if the Block PDF has been saved to disk using *File, Save* or *File, Save As...* . You can identify modified Blocks by the asterisk after the Block name. The Preview feature can be configured to save the Block PDF automatically before creating a Preview. This way you can make sure that interactive changes will immediately be reflected in the Preview.

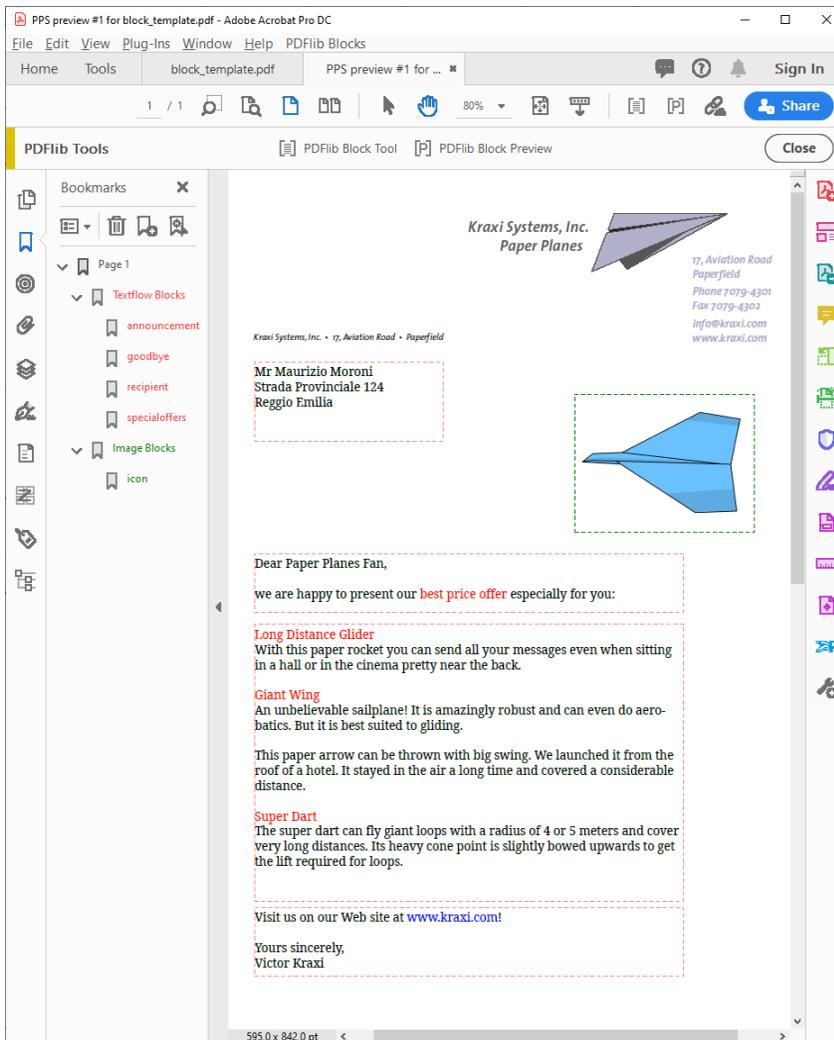


Fig. 13.4
Preview PDF for the
document shown in
figure 13.1. It contains
Block info layers and
annotations

Configuring the Preview. Several aspects of Block Preview creation and the underlying PPS operation can be configured via *PDFlib Blocks, Preview, Preview Configuration...*:

- ▶ Preview for the current page or the full document;
- ▶ Output directory for the generated Preview documents;
- ▶ Automatically save the Block PDF before creating the Preview;
- ▶ Add Block info layers and annotations;
- ▶ Copy Blocks to the generated output;
- ▶ *Clone PDF/A, PDF/UA or PDF/X status of Block PDF*: since these standards restrict the use of layers and annotations the *Block info layers and annotations* option is mutually exclusive with this option.
- ▶ *Copy Blocks to Preview File* allows you copy the PDFlib Blocks to the generated Preview upon filling. All Blocks will be copied, regardless of whether or not they could successfully be filled.

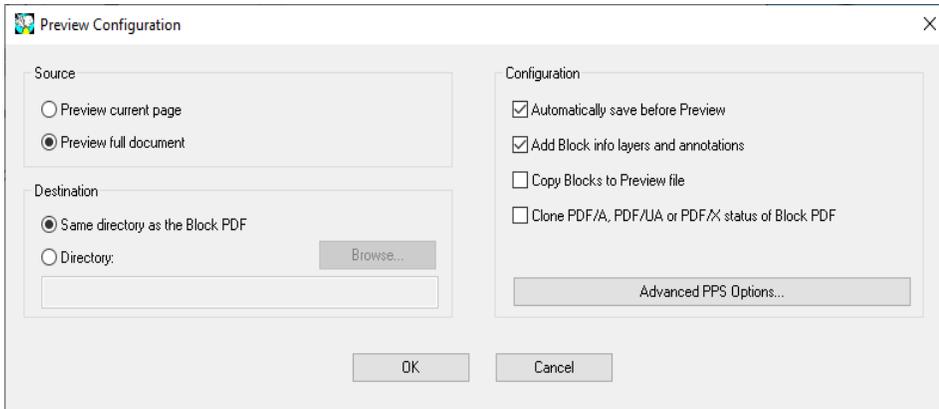


Fig. 13.5 Block Preview configuration

- ▶ The *Advanced PPS options* dialog can be used to specify additional option lists for PPS functions according to the PPS API. For example, the *searchpath* option for *PDF_set_option()* can be used to specify a directory where fonts or images for Block filling are located. It is recommended to specify advanced options in cooperation with the developer of the PPS code.

Block ordering. When the document is stored with »Save as...« in Acrobat the Blocks are sorted in alphabetical order according to the Block name. This is also the ordering in which Blocks are processed by Preview and reported by pCOS. However, applications will typically fill Blocks based on their name (as opposed to the storage order on file), so the ordering in the PDF document is usually not relevant.

Block filtering and Preview. If a Block filter (see Section 13.3.3, »Block Filtering«, page 343) is active the filter affects both screen display and Preview. Filtered Blocks are therefore not processed when creating the Preview.

Information provided with the Preview. The generated Preview documents contain the original page contents (the background), the filled Blocks, and optionally various other pieces of information. This information can be useful for checking and improving Blocks and PPS configuration. The following items will be created for each active Block with default contents:

- ▶ *Error markers:* Blocks which could not be filled successfully are visualized by a crossed-out rectangle so that they can easily be identified. Error markers will always be created if a Block couldn't be processed.
- ▶ *Bookmarks:* The processed Blocks will be summarized in bookmarks which are structured according to the page number, the Block type, and possible errors. Bookmarks can be displayed via *View, Show/Hide, Navigation Panes, Bookmarks*. Bookmarks will always be created.
- ▶ *Annotations:* For each processed Block an annotation is created on the page in addition to the actual Block contents. The annotation rectangle visualizes the original Block boundary (depending on the default contents and filling mode this may differ from the boundary of the Block contents). The annotation contains the name of the Block and an error message if the Block couldn't be filled. Annotations are generated

by default, but can be disabled in the Preview configuration. Since the use of annotations is restricted in the PDF/A and PDF/X standards, annotations are not created if the *Clone PDF/A, PDF/UA or PDF/X status of Block PDF* option is enabled.

- ▶ **Layers:** The page contents are placed on layers to facilitate analysis and debugging. A separate layer are created for the page background (i.e. the contents of the original page), each Block type, error Blocks which couldn't be filled, and the annotations with Block information. If a layer remains empty (e.g. no errors occurred) it will not be created. The layer list can be displayed via *View, Navigation Panels, Layers*. By default, all layers on the page are displayed. In order to hide the contents of a layer click on the eye symbol to the left of the layer name. Layer creation can be disabled in the Preview configuration. Since the use of layers is restricted in the standards PDF/A-1 and PDF/X-3, layers are not created if the *Clone PDF/A, PDF/UA or PDF/X status of Block PDF* option is enabled.

Cloning PDF/A, PDF/UA or PDF/X status. The *Clone PDF/A, PDF/UA or PDF/X status of Block PDF* configuration is useful when PDF output according to one of these standards must be created. Clone mode can be enabled if the input conforms to one of the following standards:

PDF/A-1a:2005, PDF/A-1b:2005
PDF/A-2a, PDF/A-2b, PDF/A-2u
PDF/A-3a, PDF/A-3b, PDF/A-3u

PDF/UA-1

PDF/X-4
PDF/X-4p
PDF/X-5n

When Previews are created in clone mode, PPS duplicates the following aspects of the Block PDF in the generated Preview:

- ▶ the PDF standard identification;
- ▶ output intent condition;
- ▶ page sizes including all page boxes;
- ▶ Tagged PDF: structure hierarchy and document language (if present);
- ▶ XMP document metadata.

When cloning standard-conforming PDF documents all Block filling operations must conform to the respective standard. For example, if no output intent is present RGB images without ICC profile can not be used. Similarly, all used fonts must be embedded. The full list of requirements can be found in Section 12.2, »PDF/A for Archiving«, page 305, and Section 12.3, »PDF/X for Print Production«, page 316. If a Block filling operation in PDF/A or PDF/X cloning mode would violate the selected standard (e.g. because a default image uses RGB color space, but the document does not contain a suitable output intent) an error message pops up and no Preview will be generated. This way users can catch potential standard violations very early in the workflow.

13.6 Filling Blocks with PPS

In order to fill Blocks with PPS you must first place the page containing the Blocks on the output page with the `PDF_fit_pdi_page()` method. After placing the page its Blocks can be filled with the `PDF_fill_*block()` functions.

Simple example: add variable text to a template. Adding dynamic text to a PDF template is a very common task. The following code fragment opens a page in an input PDF document (the template or Block container), places it on the output page, and fills some variable text into a text Block called `firstname`:

```
doc = p.open_pdi_document(filename, "");
if (doc == -1)
    throw new Exception("Error: " + p.get_errmsg());

page = p.open_pdi_page(doc, pageno, "");
if (page == -1)
    throw new Exception("Error: " + p.get_errmsg());

p.begin_page_ext(width, height, "");
/* Place the imported page */
p.fit_pdi_page(page, 0.0, 0.0, "");

/* Fill a single Block on the placed page */
p.fill_textblock(page, "firstname", "Serge", "");

p.close_pdi_page(page);
p.end_page_ext("");
p.close_pdi_document(doc);
```

Cookbook A full code sample can be found in the *Cookbook* topic `blocks/starter_block`.

Overriding Block properties. In certain situations the programmer wants to use only some of the properties provided in a Block definition, but override other properties with custom values. This can be useful in various situations:

- ▶ Business logic may decide to enforce certain overrides.
- ▶ The scaling factor for an image or PDF page will be calculated by the application instead of taken from the Block definition.
- ▶ Change the Block coordinates programmatically, for example when generating an invoice with a variable number of data items.
- ▶ Individual spot color names could be supplied in order to match customer requirements in a print shop application.

Property overrides can be achieved by supplying property names and the corresponding values in the option list of the `PDF_fill_*block()` functions, e.g.

```
p.fill_textblock(page, "firstname", "Serge", "fontsize=12");
```

This will override the Block's internal `fontsize` property with the supplied value 12. Almost all property names can be used as options.

Property overrides apply only to the respective method calls; they are not stored in the Block definition.

Moving Blocks based on the contents of previous Blocks. The fixed size of a Block may not match its variable contents. For example, if only few text is filled into a large Block a gap between two Blocks may arise; if there is too much text it may not fit into the Block rectangle. In this situation you can query the results of the Block filling process and adjust the position of subsequent Blocks:

- ▶ The default *fitmethod* is *auto*, i.e. the text is forced to fit into the Block rectangle. To allow excess contents to overflow the Block you must set *fitmethod* to *nofit*. This can be specified in the Block properties in the Block template at design time or by supplying the *fitmethod* option to `PDF_fill_textblock()`.
- ▶ Supply the option *matchbox={name=...}* to `PDF_fill_*block()` to create a named matchbox for formatted contents. The name can be chosen arbitrarily, but must be unique on the page. A simple approach is to use the Block name as matchbox name..
- ▶ Query the lower vertical position of the matchbox with `PDF_info_matchbox()` to determine the extent of the formatted Block contents. and the pCOS path `pages[.]/blocks/<blockname>/Rect[1]` with `PDF_pcos_get_number()` to determine the position of the lower edge of the Block rectangle.
- ▶ Calculate the difference between both values. If this offset is positive the contents didn't completely fill the Block; if it is negative the contents overflowed the Block. In both cases you can move the next Block up or down by the calculated offset. This can be achieved with the *refpoint* option of `PDF_fill_*block()` which overrides the *Rect* property. This option supports relative coordinates, i.e. you can supply the relative offset with the suffix *r* instead of calculating the new absolute position.
- ▶ You can apply this method to an arbitrary number of Blocks (of any type) by accumulating the per-Block offsets. Depending on the contents each Block will move successive Blocks up or down by an appropriate amount.

Cookbook A full code sample can be found in the `starter_block` sample.

Placing the imported page on top of the filled Blocks. The imported page must have been placed on the output page before using any of the Block filling functions. This means that the original page will usually be placed below the Block contents. However, in some situations it may be desirable to place the original page on top of the filled Blocks. This can be achieved by placing the page once with the *blind* option of `PDF_fit_pdi_page()` in order to make its Blocks and their position known to PPS, and place it again after filling the Blocks in order to actually show the page contents:

```
/* Place the page in blind mode to prepare the Blocks, without the page being visible */
p.fit_pdi_page(page, 0.0, 0.0, "blind");

/* Fill the Blocks */
p.fill_textblock(page, "firstname", "Serge", "");
/* ... fill more Blocks ... */

/* Place the page again, this time visible */
p.fit_pdi_page(page, 0.0, 0.0, "");
```

Cookbook A full code sample can be found in the *Cookbook* topic `blocks/block_below_contents`.

Ignoring the container page when filling Blocks. Imported Blocks can also be useful as placeholders without any reference to the underlying contents of the Block's page. You can place a container page with Blocks in blind mode on one or more pages, i.e. with the

blind option of `PDF_fit_pdi_page()`, and subsequently fill its Blocks. This way you can take advantage of Blocks and their properties without placing the container page on the output page, and can duplicate Blocks on multiple pages (or even on the same output page).

Cookbook A full code sample can be found in the *Cookbook topic* `blocks/duplicate_block`.

Linking Textflow Blocks. Textflow Blocks can be linked so that one Block holds the overflow text of a previous Block. For example, if you have long variable text which may need to be continued on another page you can link two Blocks and fill the remaining text of the first Block into the second Block.

PPS internally creates a Textflow from the text provided to `PDF_fill_textblock()` and the Block properties. For unlinked Blocks this Textflow is placed in the Block and the corresponding Textflow handle is deleted at the end of the call; overflow text is lost in this case.

With linked Textflow Blocks the overflow text of the first Block can be filled into the next Block. The remainder of the first Textflow is used as Block contents instead of creating a new Textflow. Linking Textflow Blocks works as follows:

- ▶ In the first call to `PDF_fill_textblock()` within a chain of linked Textflow Blocks the value `-1` (in PHP: `0`) must be supplied for the `textflowhandle` option. The Textflow handle created internally is returned by `PDF_fill_textblock()`, and must be stored by the application.
- ▶ In the next call to `PDF_fill_textblock()` the Textflow handle returned in the previous step can be supplied to the `textflowhandle` option (the text supplied in the `text` parameter is ignored in this case, and should be empty). The Block is filled with the remainder of the Textflow.
- ▶ This process can be repeated with more Textflow Blocks.
- ▶ The returned Textflow handle can be supplied to `PDF_info_textflow()` in order to determine the results of Block filling, e.g. the end condition or the end position of the text.

Note that the `fitmethod` property should be set to `clip` (this is the default anyway if `textflowhandle` is supplied). The basic code fragment for linking Textflow Blocks looks as follows:

```
p.fit_pdi_page(page, 0.0, 0.0, "");
tf = -1;

for (i = 0; i < blockcount; i++)
{
    String optlist = "textflowhandle=" + tf;
    int reason;
    tf = p.fill_textblock(page, blocknames[i], text, optlist);
    text = null;

    if (tf == -1)
        break;

    /* check result of most recent call to fit_textflow() */
    reason = (int) p.info_textflow(tf, "returnreason");
    result = p.get_string(reason, "");

    /* end loop if all text was placed */
}
```

```

    if (result.equals("_stop"))
    {
        p.delete_textflow(tf);
        break;
    }
}

```

Cookbook A full code sample can be found in the *Cookbook* topic `blocks/linked_textblocks`.

Block filling order. The Block functions `PDF_fill_*block()` process properties and Block contents in the following order:

- ▶ **Background:** if the `backgroundcolor` property is present and contains a color space keyword different from `None`, the Block area will be filled with the specified color.
- ▶ **Border:** if the `bordercolor` property is present and contains a color space keyword different from `None`, the Block border will be stroked with the specified color and line-width.
- ▶ **Contents:** the supplied Block contents and all other properties except `bordercolor` and `linewidth` will be processed.
- ▶ **Textline and Textflow Blocks:** if neither text nor default text has been supplied, there won't be any output at all, not even background color or Block border.

Nested Blocks. Before Blocks can be filled the page containing the Blocks must have been placed on the output page before (since otherwise PPS wouldn't know the location of the Blocks after scaling, rotating, and translating the page). If the page only serves as a Block container without bringing static content to the new page you can place the imported page with the `blind` option.

For successful Block filling it doesn't matter how the imported page was placed on the output page:

- ▶ The page can be placed directly with `PDF_fit_pdi_page()`.
- ▶ The page can be placed indirectly in a table cell with `PDF_fit_table()`.
- ▶ The page can be placed as contents of a another PDF Block with `PDF_fill_pdfblock()`.

The third method, i.e. filling a PDF Block with another page containing Blocks, allows nested Block containers. This allows simple implementations of interesting use cases. For example, you can implement both imposition and personalization with a two-step Block filling process:

- ▶ The first-level Block container page contains several large PDF Blocks which indicate the major areas on the paper to be printed on. The arrangement of PDF Blocks reflects the intended post-processing of the paper (e.g. folding or cutting).
- ▶ Each of the first-level PDF Blocks is then filled with a second-level container PDF page which contains Text, Image, PDF, or Graphics Blocks to be filled with variable text for personalization.

With this method Block containers can be nested. Although Block nesting works to an arbitrary level, a nesting level of three or more will only rarely be required.

The second-level Block containers (e.g. a template page for a letter) may be identical or different for each imposed page. If they are identical the Blocks on the letter template must be filled before placing the letter template itself in the next first-level Block since PPS always uses the location of the most recent placement of the template page.

Cookbook A full code sample can be found in the *Cookbook* topic `blocks/nested_blocks`.

Block coordinates. The Rectangle coordinates of a Block refer to the PDF default coordinate system. When the page containing the Block is placed on the output page with PPS, several positioning and scaling options can be supplied to `PDF_fit_pdi_page()`. These options are taken into account when the Block is being processed. This makes it possible to place a template page on the output page multiply, every time filling its Blocks with data. For example, a business card template may be placed four times on an imposition sheet. The Block functions will take care of the coordinate system transformations, and correctly place the text for all Blocks in all invocations of the page. The only requirement is that the client must place the page and then process all Blocks on the placed page. Then the page can be placed again at a different location on the output page, followed by more Block processing operations referring to the new position, and so on.

The Block Plugin displays the Block coordinates differently from what is stored in the PDF file. The plugin uses Acrobat's convention which has the coordinate origin in the upper left corner of the page, while the internal coordinates (those stored in the Block) use PDF's convention of having the origin at the lower left corner of the page. The coordinate display in the Properties dialog is also subject to the units specified in Acrobat (see »Block size and position«, page 338).

Spot colors in Block properties. To use a separation (spot) color in a Block property you can click the »...« button which will present a list of all HKS and Pantone spot colors. These color names are built into PPS and can be used without further preparations. For custom spot colors an alternate color can be defined in the Block Plugin. If no alternate color is specified in the Block properties, the custom spot color must have been defined earlier in the PPS application using `PDF_makespotcolor()` or a suitable color option list. Otherwise Block filling will fail.

13.7 Block Properties

PPS and the Block Plugin support general properties which can be assigned to any type of Block. In addition there are properties which are specific to the Block types *Textline*, *Textflow*, *Image*, *PDF*, and *Graphics*.

Properties support the same data types as option lists except handles and action lists. The names of Block properties are generally identical to options for API Methods such as `PDF_fit_textline()`, `PDF_fit_image()` (e.g., *fitmethod*, *charspacing*). In these cases the behavior is exactly the same as the one documented for the respective option.

13.7.1 Administrative Properties

Administrative properties apply to all Block types. Required entries will automatically be generated by the Block Plugin. Table 13.4 lists the administrative Block properties.

Table 13.4 Administrative properties

keyword	possible values and explanation
Description	(String) Human-readable description of the Block's function, encoded in PDFDocEncoding or Unicode (in the latter case starting with a BOM). This property is for user information only, and is ignored by PPS.
Locked	(Boolean) If <code>true</code> , the Block and its properties can not be edited with the Block Plugin. This property will be ignored by PPS. Default: <code>false</code>
Name	(String; required) Name of the Block. Block names must be unique within a page, but not within a document. The three characters [] / are not allowed in Block names. Block names are restricted to a maximum of 125 characters.
Subtype	(Keyword; required) Depending on the Block type, one of <code>Text</code> , <code>Image</code> , <code>PDF</code> , or <code>Graphics</code> . Note that <code>Textline</code> and <code>Textflow</code> Blocks both have Subtype <code>Text</code> , but are distinguished by the <code>textflow</code> property.
textflow	(Boolean) Controls single- or multiline processing. This property is not available explicitly in the user interface of the Block Plugin, but will be mapped to <code>Textline</code> or <code>Textflow</code> Blocks, respectively (Default: <code>false</code>): false Textline Block: text spans a single line and will be processed with <code>PDF_fit_textline()</code> . true Textflow Block: text can span multiple lines and will be processed with <code>PDF_fit_textflow()</code> . In addition to the standard text properties Textflow-related properties can be specified (see Table 13.9).
Type	(Keyword; required) Always <code>Block</code>

13.7.2 Rectangle Properties

Rectangle properties apply to all Block types. They describe the appearance of the Block rectangle itself. Required entries will automatically be generated by the Block Plugin. Table 13.5 lists the rectangle properties.

Table 13.5 Rectangle properties

keyword	possible values and explanation
background-color	(Color) If this property is present and contains a color space keyword different from None, a rectangle will be drawn and filled with the supplied color. This may be useful to cover existing page contents. Default: None
bordercolor	(Color) If this property is present and contains a color space keyword different from None, a rectangle will be drawn and stroked with the supplied color. Default: None
linewidth	(Float; must be greater than 0) Stroke width of the line used to draw the Block rectangle; only used if bordercolor is set. Default: 1
Rect	(Rectangle; required) The Block coordinates. The origin of the coordinate system is in the lower left corner of the page. However, the Block Plugin displays the coordinates in Acrobat's notation, i.e., with the origin in the upper left corner of the page. The coordinates will be displayed in the unit which is currently selected in Acrobat, but will always be stored in points in the PDF file.
Status	(Keyword) Describes how the Block will be processed by PPS and the Preview feature (default: active): active The Block will be fully processed according to its properties. ignore The Block will be ignored. ignoredefault Like active, except that the defaulttext/image/pdf/graphics properties and options are ignored, i.e. the Block remains empty if no variable contents are available (especially in the Preview). This may be useful to make sure that the Block's default contents are not used for filling Blocks on the server side although the Block may contain default contents for generating Previews. It can also be used to disable the default contents for previewing a Block without removing the default contents from the Block properties. static No variable contents will be placed; instead, the Block's default text, image, PDF, or graphics contents will be used if available.

13.7.3 Appearance Properties

Appearance properties specify formatting details:

- ▶ Table 13.6 lists transparency appearance properties for all Block types.
- ▶ Graphics state properties (see options of `PDF_create_gstate()` in the PDFlib API Reference).
- ▶ Table 13.7 lists text appearance properties for Textline and Textflow Blocks.

Table 13.6 Transparency appearance properties for all Block types

keyword	possible values and explanation
blendmode	(Keyword list; if used in PDF/A-1 mode it must have the value Normal) Name of the blend mode: None, Color, ColorDodge, ColorBurn, Darken, Difference, Exclusion, HardLight, Hue, Lighten, Luminosity, Multiply, None, Normal, Overlay, Saturation, Screen, SoftLight. <i>Default:</i> None
opacityfill	(Float; if used in PDF/A mode it must have the value 1) Opacity for fill operations in the range 0..1. The value 0 means fully transparent; 1 means fully opaque.
opacitystroke	(Float; if used in PDF/A mode it must have the value 1) Opacity for stroke operations in the range 0..1. The value 0 means fully transparent; 1 means fully opaque.

Table 13.7 Text appearance properties for Textline and Textflow Blocks

keyword	possible values and explanation																		
charspacing	(Float or percentage) Character spacing. Percentages are based on fontsize. Default: 0																		
decoration-above	(Boolean) If true, the text decoration enabled with the underline, strikethrough, and overline options will be drawn above the text, otherwise below the text. Changing the drawing order affects visibility of the decoration lines. Default: false																		
fillcolor	(Color) Fill color of the text. Default: gray 0 (=black)																		
fontname	(String; required) Name of the font. The Block plugin presents a list of fonts available in the system. However, these font names may not be portable between macOS, Windows, and Unix systems. If fontname starts with an '@' character the font is applied in vertical writing mode.																		
fontsize	(Float; required) Size of the font in points																		
horizscaling	(Float or percentage) Horizontal text scaling. Default: 100%																		
italicangle	(Float) Italic angle of text in degrees. Default: 0																		
kerning	(Boolean) Kerning behavior. Default: false																		
overline	(Boolean) Overline mode. Default: false																		
shadow	(Composite) Create a shadow effect (default: no shadow). The following subproperties are available: fillcolor (Color) Color of the shadow. Default: {gray 0.8} offset (List of 2 floats or percentages) The shadow's offset from the reference point of the text in user coordinates or as a percentage of the font size. Default: {5% -5%}																		
strikethrough	(Boolean) Strikethrough mode. Default: false																		
strokecolor	(Color) Stroke color of the text. Default: gray 0 (=black)																		
strokewidth	(Float, percentage, or keyword; only effective if textrendering is set to stroke text) Line width for outline text (in user coordinates or as a percentage of the fontsize). The keyword auto or the equivalent value o uses a built-in default. Default: auto																		
textrendering	(Integer) Text rendering mode. Only the value 3 has an effect on Type 3 fonts (default: 0): <table border="0" style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 50%; vertical-align: top;"> <table border="0"> <tr> <td style="width: 20px;">0</td> <td> fill text</td> </tr> <tr> <td>1</td> <td> stroke text (outline)</td> </tr> <tr> <td>2</td> <td> fill and stroke text</td> </tr> <tr> <td>3</td> <td>invisible text</td> </tr> </table> </td> <td style="width: 50%; vertical-align: top;"> <table border="0"> <tr> <td style="width: 20px;">4</td> <td> fill text and add it to the clipping path</td> </tr> <tr> <td>5</td> <td> stroke text and add it to the clipping path</td> </tr> <tr> <td>6</td> <td> fill and stroke text and add it to the clipping path</td> </tr> <tr> <td>7</td> <td> add text to the clipping path (not for Blocks)</td> </tr> </table> </td> </tr> </table>	<table border="0"> <tr> <td style="width: 20px;">0</td> <td> fill text</td> </tr> <tr> <td>1</td> <td> stroke text (outline)</td> </tr> <tr> <td>2</td> <td> fill and stroke text</td> </tr> <tr> <td>3</td> <td>invisible text</td> </tr> </table>	0	 fill text	1	 stroke text (outline)	2	 fill and stroke text	3	invisible text	<table border="0"> <tr> <td style="width: 20px;">4</td> <td> fill text and add it to the clipping path</td> </tr> <tr> <td>5</td> <td> stroke text and add it to the clipping path</td> </tr> <tr> <td>6</td> <td> fill and stroke text and add it to the clipping path</td> </tr> <tr> <td>7</td> <td> add text to the clipping path (not for Blocks)</td> </tr> </table>	4	 fill text and add it to the clipping path	5	 stroke text and add it to the clipping path	6	 fill and stroke text and add it to the clipping path	7	 add text to the clipping path (not for Blocks)
<table border="0"> <tr> <td style="width: 20px;">0</td> <td> fill text</td> </tr> <tr> <td>1</td> <td> stroke text (outline)</td> </tr> <tr> <td>2</td> <td> fill and stroke text</td> </tr> <tr> <td>3</td> <td>invisible text</td> </tr> </table>	0	 fill text	1	 stroke text (outline)	2	 fill and stroke text	3	invisible text	<table border="0"> <tr> <td style="width: 20px;">4</td> <td> fill text and add it to the clipping path</td> </tr> <tr> <td>5</td> <td> stroke text and add it to the clipping path</td> </tr> <tr> <td>6</td> <td> fill and stroke text and add it to the clipping path</td> </tr> <tr> <td>7</td> <td> add text to the clipping path (not for Blocks)</td> </tr> </table>	4	 fill text and add it to the clipping path	5	 stroke text and add it to the clipping path	6	 fill and stroke text and add it to the clipping path	7	 add text to the clipping path (not for Blocks)		
0	 fill text																		
1	 stroke text (outline)																		
2	 fill and stroke text																		
3	invisible text																		
4	 fill text and add it to the clipping path																		
5	 stroke text and add it to the clipping path																		
6	 fill and stroke text and add it to the clipping path																		
7	 add text to the clipping path (not for Blocks)																		
textrise	(Float or percentage) Text rise parameter. Percentages are based on fontsize. Default: 0																		
underline	(Boolean) Underline mode. Default: false																		
underline-position	(Float, percentage, or keyword) Position of the stroked line for underlined text relative to the baseline. Percentages are based on fontsize. Default: auto																		
underline-width	(Float, percentage, or keyword) Line width for underlined text. Percentages are based on fontsize. Default: auto																		
wordspacing	(Float or percentage) Word spacing. Percentages are based on fontsize. Default: 0																		

13.7.4 Text Preparation Properties

Text preparation properties specify preprocessing steps for Textline and Textflow Blocks. Table 13.8 lists text preparation properties for Textline and Textflow Blocks.

Table 13.8 Text preparation properties for Textline and Textflow Blocks

keyword	possible values and explanation
charref	(Boolean) If <code>true</code> , enable substitution of numeric and character entity references and glyph name references. Default: the global <code>charref</code> option
escape-sequence	(Boolean) If <code>true</code> , enable substitution of escape sequences in content strings, hypertext strings, and name strings. Default: the global <code>escapesequence</code> option
features	<p>(List of keywords) Specifies which typographic features of an OpenType font are applied to the text, subject to the <code>script</code> and <code>language</code> options. Features which are not present in the font are ignored:</p> <p>_none Apply none of the features in the font.</p> <p><name> Enable a feature by supplying its four-character OpenType name. Common feature names are <code>aalt</code>, <code>liga</code>, <code>ital</code>, <code>tnum</code>, <code>smcp</code>, <code>swsh</code>, <code>zero</code>. A list with names and descriptions of all supported features can be found in Section 7.4, »OpenType Layout Features«, page 139</p> <p>no<names> The prefix <code>no</code> in front of a feature name (e.g. <code>noliga</code>) disables this feature.</p> <p>The following features are enabled by default: <code>calt</code>, <code>ccmp</code>, <code>clig</code>, <code>liga</code>, <code>locl</code>. The features <code>vert2</code> and <code>vert</code> are automatically enabled for fonts in vertical writing mode.</p>
language	(Keyword; only relevant if <code>script</code> is supplied) The text will be processed according to the specified language, which is relevant for the <code>features</code> and <code>shaping</code> options. A full list of keywords can be found in Section 7.5.2, »Controlling the Shaping Process«, page 146, e.g. <code>ARA</code> (Arabic), <code>JAN</code> (Japanese), <code>HIN</code> (Hindi). Default: <code>_none</code> (undefined language)
script	<p>(Keyword; required if <code>shaping=true</code>) The text is processed according to the specified script which is relevant for the options <code>features</code>, <code>shaping</code>, and <code>advancedlinebreak</code> (default: <code>_none</code>):</p> <p>_auto (Only relevant for <code>shaping</code>) Selects the script to which the majority text belongs, where <code>latn</code> and <code>_none</code> are ignored.</p> <p>_none Undefined script.</p> <p>The following scripts are supported for complex script shaping:</p> <p>arab, hebr Middle East scripts</p> <p>sinh, tibt Central Asian scripts</p> <p>khmr, thai Southeast Asian scripts</p> <p>hang, hani East Asian scripts</p> <p>beng, deva, gujr, knda, mlym, orya, guru, taml, telu Indic scripts</p> <p>Additional script tags according to the OpenType specification, e.g. <code>latn</code>, <code>cyrl</code>, <code>grek</code>, can be used to control OpenType feature processing. A full list of keywords can be found in Section 7.5.2, »Controlling the Shaping Process«, page 146. Default: <code>_none</code></p>
shaping	(Boolean) If <code>true</code> , complex script shaping and bidirectional reordering are applied to the text according to the <code>script</code> and <code>language</code> options. The <code>script</code> option must have a value different from <code>_none</code> and the font must obey certain conditions. Shaping is not available for right-to-left text in Textflows (only in Textlines). Default: <code>false</code>

13.7.5 Text Formatting Properties

Table 13.9 lists properties which can only be used for Textflow Blocks, with the exception of the *stamp* property which can also be used for Textline Blocks. They will be used to construct the initial option list for processing the Textflow (corresponding to the *optlist* parameter of `PDF_create_textflow()`). Inline option lists for Textflows can not be specified with the plugin, but they can be supplied on the server as part of the text contents when filling the Block with `PDF_fill_textblock()`, or in the Block's *defaulttext* property.

Table 13.9 Text formatting properties (mostly for Textflow Blocks)

keyword	possible values and explanation
adjust-method	(Keyword) Method used to adjust a line when a text portion doesn't fit into a line after compressing or expanding the distance between words subject to the limits specified by the <i>minspacing</i> and <i>maxspacing</i> options (default: <i>auto</i>): auto The following methods are applied in order: <i>shrink</i> , <i>spread</i> , <i>nofit</i> , <i>split</i> . clip Same as <i>nofit</i> , except that the long part at the right edge of the fit box (taking into account the <i>rightindent</i> option) will be clipped. nofit The last word will be moved to the next line provided the remaining (short) line will not be shorter than the percentage specified in the <i>nofitlimit</i> option. Even justified paragraphs may look slightly ragged. shrink If a word doesn't fit in the line the text will be compressed subject to <i>shrinklimit</i> . If it still doesn't fit the <i>nofit</i> method will be applied. split The last word will not be moved to the next line, but will forcefully be hyphenated. For text fonts a hyphen character will be inserted, but not for symbol fonts. spread The last word will be moved to the next line and the remaining (short) line will be justified by increasing the distance between characters in a word, subject to <i>spreadlimit</i> . If justification still cannot be achieved the <i>nofit</i> method will be applied.
advanced-linebreak	(Boolean) Enable the advanced line breaking algorithm which is required for complex scripts. This is required for linebreaking in scripts which do not use space characters for designating word boundaries, e.g. Thai. The options <i>locale</i> and <i>script</i> will be honored. Default: <i>false</i>
alignment	(Keyword) Specifies formatting for lines in a paragraph. Default: <i>left</i> . left left-aligned, starting at <i>leftindent</i> center centered between <i>leftindent</i> and <i>rightindent</i> right right-aligned, ending at <i>rightindent</i> justify left- and right-aligned
avoid-emptybegin	(Boolean) If <i>true</i> , empty lines at the beginning of a fitbox will be deleted. Default: <i>false</i>
fixedleading	(Boolean) If <i>true</i> , the first leading value found in each line will be used. Otherwise the maximum of all leading values in the line will be used. Default: <i>false</i>
hortab-method	(Keyword) Treatment of horizontal tabs in the text. If the calculated position is to the left of the current text position, the tab will be ignored (default: <i>relative</i>): relative The position will be advanced by the amount specified in <i>hortabsize</i> . typewriter The position will be advanced to the next multiple of <i>hortabsize</i> . ruler The position will be advanced to the <i>n</i> -th tab value in the <i>ruler</i> option, where <i>n</i> is the number of tabs found in the line so far. If <i>n</i> is larger than the number of tab positions the <i>relative</i> method will be applied.
hortabsize	(Float or percentage) Width of a horizontal tab ¹ . The interpretation depends on the <i>hortabmethod</i> option. Default: 7.5%

Table 13.9 Text formatting properties (mostly for Textflow Blocks)

keyword	possible values and explanation
lastalignment	(Keyword) Formatting for the last line in a paragraph. All keywords of the alignment option are supported, plus the following (default: auto): auto Use the value of the alignment option unless it is justify. In the latter case left will be used.
leading	(Float or percentage) Distance between adjacent text baselines in user coordinates, or as a percentage of the font size. Default: 100%
locale	(Keyword) The locale which will be used for localized linebreaking methods if advancedlinebreak=true. The keywords consists of one or more components, where the optional components are separated by an underscore character '_' (the syntax slightly differs from NLS/POSIX locale IDs): <ul style="list-style-type: none"> ▶ A required two- or three-letter lowercase language code according to ISO 639-2, e.g. en, (English), de (German), ja (Japanese). This differs from the language option. ▶ An optional four-letter script code according to ISO 15924, e.g. Hira (Hiragana), Hebr (Hebrew), Arab (Arabic), Thai (Thai). ▶ An optional two-letter uppercase country code according to ISO 3166, e.g. DE (Germany), CH (Switzerland), GB (United Kingdom) Specifying a locale is not required for advanced line breaking: the keyword <code>_none</code> specifies that no locale-specific processing will be done. Default: <code>_none</code> Examples: <code>de_DE</code> , <code>en_US</code> , <code>en_GB</code>
maxspacing minspacing	(Float or percentage) The maximum or minimum distance between words (in user coordinates, or as a percentage of the width of the space character). The calculated word spacing is limited by the provided values (but the wordspacing option will still be added). Defaults: <code>minspacing=50%</code> , <code>maxspacing=500%</code>
minlinecount	(Integer) Minimum number of lines in the last paragraph of the fitbox. If there are fewer lines they will be placed in the next fitbox. The value 2 can be used to prevent single lines of a paragraph at the end of a fitbox («orphans»). Default: 1
nofitlimit	(Float or percentage) Lower limit for the length of a line with the nofit method (in user coordinates or as a percentage of the width of the fitbox). Default: 75%
parindent	(Float or percentage) Left indent of the first line of a paragraph ¹ . The amount will be added to leftindent. Specifying this option within a line will act like a tab. Default: 0
rightindent leftindent	(Float or percentage) Right or left indent of all text lines ¹ . If leftindent is specified within a line and the determined position is to the left of the current text position, this option will be ignored for the current line. Default: 0
ruler²	(List of floats or percentages) List of absolute tab positions for hortabmethod=ruler ¹ . The list may contain up to 32 non-negative entries in ascending order. Default: integer multiples of hortabsize
shrinklimit	(Percentage) Lower limit for compressing text with the shrink method; the calculated shrinking factor is limited by the provided value, but will be multiplied with the value of the horiszscaling option. Default: 85%
spreadlimit	(Float or percentage) Upper limit for the distance between two characters for the spread method (in user coordinates or as a percentage of the font size); the calculated character distance will be added to the value of the charspacing option. Default: 0
stamp	(Keyword; Textline and Textflow Blocks) This option can be used to create a diagonal stamp within the Block rectangle. The text comprising the stamp will be as large as possible. The options position, fitmethod, and orientate (only north and south) will be honored when placing the stamp text in the box. Default: none. llzur The stamp will run diagonally from the lower left corner to the upper right corner. ulzlr The stamp will run diagonally from the upper left corner to the lower right corner. none No stamp will be created.

Table 13.9 Text formatting properties (mostly for Textflow Blocks)

keyword	possible values and explanation
tabalignchar	(Unichar) Unicode value of the character at which decimal tabs will be aligned. Default: the period character '.' (U+002E)
tabalignment²	(List of keywords) Alignment for tab stops. Each entry in the list defines the alignment for the corresponding entry in the ruler option (default: left): center Text will be centered at the tab position. decimal The first instance of tabalignchar will be left-aligned at the tab position. If no tabalignchar is found, right alignment will be used instead. left Text will be left-aligned at the tab position. right Text will be right-aligned at the tab position.

1. In user coordinates, or as a percentage of the width of the fit box

2. Tab settings can be edited in the property subgroup Ruler Tabs for hortabmethod=ruler in the Block properties dialog.

13.7.6 Object Fitting Properties

Fitting properties are available for all Block types, although some properties are specific to a certain Block type. They control how the contents will be placed in the Block:

- ▶ Table 13.10 lists fitting properties for Textline, Image, PDF, and Graphics Blocks
- ▶ Table 13.11 lists fitting properties for Textflow Blocks (mostly related to aspects of vertical fitting).

The object fitting algorithm uses the Block rectangle as fitbox. Except for *fitmethod=clip* there is no clipping; if you want to make sure that the Block contents do not exceed the Block rectangle avoid *fitmethod=nofit*.

Table 13.10 Fitting properties for Textline, Image, PDF, and Graphics Blocks

keyword	possible values and explanation
alignchar	(Unichar or keyword; only for Textline Blocks) If the specified character is found in the text, its lower left corner will be aligned at the lower left corner of the Block rectangle. For horizontal text with orientate=north or south the first value supplied in the position option defines the position. For horizontal text with orientate=west or east the second value supplied in the position option defines the position. This option will be ignored if the specified alignment character is not present in the text. The value 0 and the keyword none suppress alignment characters. The specified fitmethod will be applied, although the text cannot be placed within the Block rectangle because of the forced positioning of alignchar. Default: none
dpi	(Float list; only for image Blocks) One or two values specifying the desired image resolution in pixels per inch in horizontal and vertical direction. With the value 0 the image's internal resolution will be used if available, or 72 dpi otherwise. This property will be ignored if the fitmethod property has been supplied with one of the keywords auto, meet, slice, or entire. Default: 0
fitmethod	(Keyword) Strategy to use if the supplied content doesn't fit into the Block rectangle: auto, clip, entire, meet, nofit or slice (default: meet)..
margin	(Float list; only for Textline Blocks) One or two float values describing additional horizontal and vertical reduction of the Block rectangle. Default: 0
orientate	(Keyword) Specifies the desired orientation of the content when it is placed. Possible values are north, east, south, west. Default: north
position	(Float list) One or two values specifying the position of the reference point within the content. The position is specified as a percentage within the Block. Only for Textline Blocks: the keyword auto can be used for the first value in the list. It indicates right if the writing direction of the text is from right to left (e.g. for Arabic and Hebrew text), and left otherwise (e.g. for Latin text). Default: {0 0}, i.e. the lower left corner
rotate	(Float) Rotation angle in degrees by which the Block will be rotated counter-clockwise before processing begins. The reference point is center of the rotation. Default: 0
scale	(Float list; only for image, PDF, and Graphics Blocks) One or two values specifying the desired scaling factor(s) in horizontal and vertical direction. This option will be ignored if the fitmethod property has been supplied with one of the keywords auto, meet, slice, or entire. Default: 1
shrinklimit	(Float or percentage; only for Textline Blocks) The lower limit of the shrinkage factor which will be applied to fit text with fitmethod=auto. Default: 0.75

Table 13.11 Fitting properties for Textflow Blocks

keyword	possible values and explanation
firstlinedist	<p>(Float, percentage, or keyword) The distance between the top of the Block rectangle and the baseline for the first line of text, specified in user coordinates, as a percentage of the relevant font size (the first font size in the line if <code>fixedleading=true</code>, and the maximum of all font sizes in the line otherwise), or as a keyword (default: <code>leading</code>):</p> <p>leading The leading value determined for the first line; typical diacritical characters such as À will touch the top of the fitbox.</p> <p>ascender The ascender value determined for the first line; typical characters with larger ascenders, such as <i>d</i> and <i>h</i> will touch the top of the fitbox.</p> <p>capheight The capheight value determined for the first line; typical capital uppercase characters such as <i>H</i> will touch the top of the fitbox.</p> <p>xheight The xheight value determined for the first line; typical lowercase characters such as <i>x</i> will touch the top of the fitbox.</p> <p>If <code>fixedleading=false</code> the maximum of all <code>leading</code>, <code>ascender</code>, <code>xheight</code>, or <code>capheight</code> values found in the first line will be used.</p>
fitmethod	<p>(Keyword) Strategy to use if the Block is too small for the Textflow:</p> <p>auto <code>fontsize</code> and <code>leading</code> are decreased until the text fits.</p> <p>clip Text is clipped at the Block margin (useful for linking Textflow Blocks).</p> <p>nofit Text runs beyond the bottom margin of the Block (useful for moving Blocks).</p> <p>Default: <code>clip</code> if the <code>textflowhandle</code> option is supplied, otherwise <code>auto</code></p>
lastlinedist	<p>(Float, percentage, or keyword) Will be ignored for <code>fitmethod=nofit</code> The minimum distance between the baseline for the last line of text and the bottom of the fitbox, specified in user coordinates, as a percentage of the font size (the first font size in the line if <code>fixedleading= true</code>, and the maximum of all font sizes in the line otherwise), or as a keyword. Default: <code>0</code>, i.e. the bottom of the fitbox will be used as baseline, and typical descenders will extend below the Block rectangle.</p> <p>descender The descender value determined for the last line; typical characters with descenders, such as <i>g</i> and <i>j</i> will touch the bottom of the fitbox.</p> <p>If <code>fixedleading=false</code> the maximum of all descender values found in the last line will be used.</p>
linespread-limit	<p>(Float or percentage; only for <code>verticalalign=justify</code>) Maximum amount in user coordinates or as percentage of the leading for increasing the leading for vertical justification. Default: <code>200%</code></p>
maxlines	<p>(Integer or keyword) The maximum number of lines in the fitbox, or the keyword <code>auto</code> which means that as many lines as possible will be placed in the fitbox. When the maximum number of lines has been placed <code>PDF_fit_textflow()</code> will return the string <code>_boxfull</code>.</p>
minfontsize	<p>(Float or percentage) Minimum allowed font size when text is scaled down to fit into the Block rectangle with <code>fitmethod=auto</code> when <code>shrinklimit</code> is exceeded. The limit is specified in user coordinates or as a percentage of the height of the Block. If the limit is reached the text will be created with the specified <code>minfontsize</code> as <code>fontsize</code>. Default: <code>0.1%</code></p>
orientate	<p>(Keyword) Specifies the desired orientation of the text when it is placed. Possible values are <code>north</code>, <code>east</code>, <code>south</code>, <code>west</code>. Default: <code>north</code></p>
rotate	<p>(Float) Rotate the coordinate system, using the lower left corner of the fitbox as center and the specified value as rotation angle in degrees. This results in the box and the text being rotated. The rotation will be reset when the text has been placed. Default: <code>0</code></p>

Table 13.11 Fitting properties for Textflow Blocks

keyword	possible values and explanation
verticalalign	(Keyword) Vertical alignment of the text in the fitbox (default: top):
top	Formatting will start at the first line, and continue downwards. If the text doesn't fill the fitbox there may be whitespace below the text.
center	The text will be vertically centered in the fitbox. If the text doesn't fill the fitbox there may be whitespace both above and below the text.
bottom	Formatting will start at the last line, and continue upwards. If the text doesn't fill the fitbox there may be whitespace above the text.
justify	The text will be aligned with top and bottom of the fitbox. In order to achieve this the leading will be increased up to the limit specified by <code>linespreadlimit</code> . The height of the first line will only be increased if <code>firstlinedist=leading</code> .

13.7.7 Tagged PDF Properties

Tagged PDF properties are available for all Block types. Table 13.12 lists tagging properties.

Table 13.12 Tagged PDF properties (PPS interface 11)

keyword	possible values and explanation
tag	(Option list) Options for structure element tagging according to Table 14.2. If a structure element name is provided as property, inline tags in Textflow are ignored.

13.7.8 Properties for default Contents

Properties for default contents specify how to fill the Block if no specific contents are provided. They are especially useful for the Preview feature since it will fill the Blocks with their default contents. Table 13.13 lists properties for default contents.

Table 13.13 Properties for default contents

keyword	possible values and explanation
default-graphics	(String; only for graphics Blocks) Path name of a graphics file which will be used if no graphics is supplied by the client application. ¹
defaultimage	(String; only for image Blocks) Path name of an image which will be used if no image is supplied by the client application. ¹
defaultpdf	(String; only for PDF Blocks) Path name of a PDF document which will be used if no substitution PDF is supplied by the client application. ¹
default-pdfpage	(Integer; only for PDF Blocks) Page number of the page in the default PDF document. Default: 1
defaulttext	(String; only for Textline and Textflow Blocks) Text which will be used if no variable text is supplied by the client application ²

1. It is recommended to use file names without absolute paths, and use the SearchPath feature in the PPS client application. This makes Block processing independent from platform and file system details.

2. The text will be interpreted in winansi encoding or Unicode.

13.7.9 Custom Properties

Custom properties apply to Blocks of any type of Block, and will be ignored by PPS and the Preview feature. Table 13.14 lists the naming rules for custom properties.

Table 13.14 Custom Block properties for all Block types

keyword	possible values and explanation
any name not containing the three characters [] /	(String, name, float, or float list) The interpretation of the values of custom properties is completely up to the client application; they will be ignored by PPS.

13.8 Querying Block Names and Properties with pCOS

In addition to automatic Block processing with PPS, the integrated pCOS facility can be used to enumerate Block names and query standard or custom properties.

Cookbook A full code sample for querying the properties of Blocks contained in an imported PDF can be found in the Cookbook topic `blocks/query_block_properties`.

Finding the number and names of Blocks. The client code must not even know the names or number of Blocks on an imported page since these can also be queried. The following statement returns the number of Blocks on page with number *pagenum*:

```
blockcount = (int) p.pcos_get_number(doc, "length:pages[" + pagenum + "]/blocks");
```

The following statement returns the name of Block number *blocknum* on page *pagenum* (Block and page counting start at 0):

```
blockname = p.pcos_get_string(doc,
    "pages[" + pagenum + "]/blocks[" + blocknum + "]/Name");
```

The returned Block name can subsequently be used to query the Block's properties or fill the Block with text, image, PDF or graphics contents. If the specified Block doesn't exist an exception will be thrown. You can avoid this by using the *length* prefix to determine the number of Blocks and therefore the maximum index in the *blocks* array (keep in mind that the Block count will be one higher than the highest possible index since array indexing starts at 0).

Checking for the presence of a Block. For additional flexibility of the client application code you can check whether for the presence of a Block before attempting to fill it. This way the designer can move Blocks between pages without breaking the application which fills the Blocks.

The following code checks whether a Block with the name *foo* is present on a page:

```
/* pCOS object type "dictionary" means that the Block is present */
if (pcos_get_string(doc, "type:pages[" + pagenum + "]/blocks/" + "foo").equals("dict"))
{
    /* Block "foo" is present on the page */
}
```

Addressing Blocks by number or name. In the pCOS path syntax for addressing Block properties the following expressions are equivalent, assuming that the Block with number 6 has its *Name* property set to *foo*:

```
pages[...]/blocks[6]
pages[...]/blocks/foo
```

Querying Block coordinates. The two coordinate pairs (*llx*, *lly*) and (*urx*, *ury*) describing the lower left and upper right corner of a Block named *foo* can be queried as follows:

```
llx = p.pcos_get_number(doc, "pages[" + pagenum + "]/blocks/foo/rect[0]");
lly = p.pcos_get_number(doc, "pages[" + pagenum + "]/blocks/foo/rect[1]");
urx = p.pcos_get_number(doc, "pages[" + pagenum + "]/blocks/foo/rect[2]");
ury = p.pcos_get_number(doc, "pages[" + pagenum + "]/blocks/foo/rect[3]");
```

Note that these coordinates are provided in the default coordinate system (with the origin in the bottom left corner, possibly modified by the page's *CropBox*), while the Block Plugin displays the coordinates according to Acrobat's user interface coordinate system with an origin in the upper left corner of the page. The values queried with the pCOS pseudo object *rect* (all lowercase) take into account any relevant *CropBox/MediaBox* and *Rotate* entries and normalize the order of the coordinates. In contrast, the values queried with the native PDF key *Rect* cannot be directly used as new coordinates if a *CropBox* is present.

Note that the *topdown* option is not taken into account when querying Block coordinates.

Querying custom properties. Custom properties can be queried as in the following example, where the property *zipcode* is queried from a Block named *b1* on page *pagenum*:

```
zip = p.pcos_get_string(doc, "pages[" + pagenum + "]/blocks/b1/Custom/zipcode");
```

If you don't know which custom properties are actually present in a Block, you can determine the names at runtime. In order to find the name of the first custom property in a Block named *b1* use the following:

```
propname = p.pcos_get_string(doc, "pages[" + pagenum + "]/blocks/b1/Custom[0].key");
```

Use increasing indexes instead of *o* in order to determine the names of all custom properties. Use the *length* prefix to determine the number of custom properties.

Non-existing Block properties and default values. Use the *type* prefix to determine whether a Block or property is actually present. If the type for a path is *o* or *null* the respective object is not present in the PDF document. Note that for predefined properties this means that the default value of the property will be used.

Name space for custom properties. In order to avoid confusion when PDF documents from different sources are exchanged, it is recommended to use an Internet domain name as a company-specific prefix in all custom property names, followed by a colon ':' and the actual property name. For example, ACME corporation would use the following property names:

```
acme.com:digits  
acme.com:refnumber
```

Since standard and custom properties are stored differently in the Block, standard PPS property names (as defined in Section 13.7, »Block Properties«, page 357) will never conflict with custom property names.

13.9 Creating and Importing Blocks programmatically

13.9.1 Creating PDFlib Blocks with POCA

PDFlib Blocks can be created programmatically with the POCA interface which is included in PPS. Using POCA the required PDF data structures for Block can be prepared and then supplied to the *blocks* option of *PDF_begin/end_page_ext()*. When creating the Block definitions the requirements in Section 13.10, »PDFlib Block Specification«, page 373, must be obeyed. The Block properties must be created according to the data types listed in Section 13.7, »Block Properties«, page 357.

Cookbook A code sample for creating PDFlib Blocks with PPS can be found in the category *b1ocks* of the *PDFlib Cookbook*.

The PDFlib Block specification contains an unfortunate redundancy in that the name of a Block is recorded twice: once in the main *Blocks* dictionary of a page, and again in the *Name* entry within a particular Block dictionary. These two names must be identical in order to avoid problems when filling the Block with PPS or previewing the Block with the Block Plugin. *PDF_begin/end_page_ext()* will therefore throw an exception if the dictionary provided with the *blocks* option contains a block definition which violates the »same block name« rule. The corresponding pairs are highlighted in blue in the code sample below.

The following code fragment demonstrates the use of POCA functions for creating the Block definition shown in Section , »Block dictionary keys«, page 375:

```

/* Create the Block dictionary */
blockdict = p.poca_new("containertype=dict usage=blocks");

/* -----
 * Create a Text Block
 * -----
 */
textblock = p.poca_new("containertype=dict usage=blocks type=name key=Type value=Block");

container1 = p.poca_new("containertype=array usage=blocks " +
    "type=integer values={70 640 300 700}");

p.poca_insert(textblock, "type=array key=Rect value=" + container1);
p.poca_insert(textblock, "type=name key=Name value=job_title");
p.poca_insert(textblock, "type=name key=Subtype value=Text");
p.poca_insert(textblock, "type=name key=fitmethod value=auto");
p.poca_insert(textblock, "type=string key=fontname value=Helvetica");
p.poca_insert(textblock, "type=float key=fontsize value=12");

/* Hook up the Block in the page's Block dictionary */
p.poca_insert(blockdict, "type=dict key=job_title direct=false value=" + textblock);

/* -----
 * Create an Image Block
 * -----
 */
imageblock = p.poca_new("containertype=dict usage=blocks " +
    "type=name key=Type value=Block");

container2 = p.poca_new("containertype=array usage=blocks " +
    "type=integer values={70 440 300 600}");

p.poca_insert(imageblock, "type=array key=Rect value=" + container2);
p.poca_insert(imageblock, "type=name key=Name value=logo");
p.poca_insert(imageblock, "type=name key=Subtype value=Image");
p.poca_insert(imageblock, "type=name key=fitmethod value=auto");

/* Hook up the Block in the page's Block dictionary */
p.poca_insert(blockdict, "type=dict key=logo direct=false value=" + imageblock);

/* -----
 * Create a PDF Block
 * -----
 */
pdfblock = p.poca_new("containertype=dict usage=blocks " +
    "type=name key=Type value=Block");

container3 = p.poca_new("containertype=array usage=blocks " +
    "type=integer values={70 240 300 400}");

p.poca_insert(pdfblock, "type=array key=Rect value=" + container3);
p.poca_insert(pdfblock, "type=name key=Name value=pdflogo");
p.poca_insert(pdfblock, "type=name key=Subtype value=PDF");
p.poca_insert(pdfblock, "type=name key=fitmethod value=meet");

/* Hook up the Block in the page's Block dictionary */
p.poca_insert(blockdict, "type=dict key=pdflogo direct=false " + "value=" + pdfblock);

```

```

/* -----
 * Hook up the Block dictionary in the current page
 * -----
 */
p.end_page_ext("blocks=" + blockdict);

/* Clean up */
p.poca_delete(blockdict, "recursive");

```

13.9.2 Importing PDFlib Blocks

You can copy one or more PDFlib Blocks from the input document to the current output page with `PDF_process_pdi()` and `action=copyallblocks` or `action=copyblock` as follows:

```

if (p.process_pdi(p, doc, 0, "action=copyallblocks block={pagenumber=1}") != 1)
{
    /* Error */
}

```

This way you can implement multi-level Block filling workflows. Keep in mind that Block names must be unique on each page, i.e. you cannot import multiple Blocks with the same name to the same page. Use the `outputblockname` suboption to rename Blocks upon copying.

13.10 PDFlib Block Specification

The Block syntax conforms to the PDF Reference which specifies an extension mechanism that allows applications to store private data attached to the data structures comprising a PDF page. A description of the PDFlib Block syntax is provided here. Users who create Blocks with the Block Plugin or PDFlib don't need this information.

PDF object structure for PDFlib Blocks. The page dictionary contains a `PieceInfo` entry which has another dictionary as value. The page dictionary should also contain the key `LastModified` which contains a time stamp for the creation or last modification of the Block structures. This dictionary contains the key `PDFlib` with an application data dictionary as value. The application data dictionary contains two standard keys listed in Table 13.15.

Table 13.15 Entries in a PDFlib application data dictionary

key	value
LastModified	(Data string; required) The date and time when the Blocks on the page were created or most recently modified. This entry will be created by PDFlib when creating Blocks with the POCA interface.
Private	(Dictionary; required) A Block list (see Table 13.16)

A Block list is a dictionary containing general information about Block processing, plus a list of all Blocks on the page. Table 13.16 lists the keys in a Block list dictionary.

Table 13.16 Entries in a Block list dictionary

key	value
Blocks	(Dictionary; required) Each key is a name object containing the name of a Block; the corresponding value is the Block dictionary for this Block (see Table 13.18). The value of the Name key in the Block dictionary must be identical to the Block's name in this dictionary.
BlockProducer¹	(String) Name of the software used to create the Blocks programmatically. This entry is created by PDFlib when creating Blocks with the POCA interface.
PluginVersion¹	(String) A string containing a version identification of the Block plugin used to create the Blocks.
pdfmark¹	(Boolean) Must be true if the Block list has been generated by use of pdfmarks.
Version	(Number; required) The version number of the Block specification to which the file complies. This document describes version 11 of the Block specification.

1. Exactly one of the keys BlockProducer, PluginVersion and pdfmark must be present.

Data types for Block properties. Properties support the same data types as option lists except handles and specialized lists such as action lists. Table 13.17 details how these types are mapped to PDF data types.

Table 13.17 Data types for Block properties

Data type	PDF type and remarks
boolean	(Boolean)
string	(String)
keyword (name)	(Name) It is an error to provide keywords outside the list of keywords supported by a particular property.
float, integer	(Number) While option lists support point and comma as decimal separators, PDF numbers require point.
percentage	(Array with two elements) The first element in the array is the number, the second element is a string containing a percent character.
list	(Array)
color	(Array with two or three elements) The first element in the array specifies a color space, and the second element specifies a color value. To specify the absence of color the respective property must be omitted. The following entries are supported for the first element in the array: /DeviceGray The second element is a single gray value. /DeviceRGB The second element is an array of three RGB values. /DeviceCMYK The second element is an array of four CMYK values. [/Separation/spotname] The first element is an array containing the keyword Separation and a spot color name. The second element is a tint value. The optional third element in the array specifies an alternate color for the spot color, which is itself a color array in one of the DeviceGray, DeviceRGB, DeviceCMYK, or Lab color spaces. If the alternate color is missing, the spot color name must either refer to a color which is known internally to PPS, or which has been defined by the application at runtime. [/Lab] The first element is an array containing the keyword Lab. The second element is an array of three Lab values.
unicar	(Text string) Unicode string in utf16be format, starting with the BOM U+FEFF

Block dictionary keys. Block dictionaries may contain the keys in Table 13.18.

Table 13.18 Entries in a Block dictionary

property group	values
<i>administrative properties</i>	<i>(Some keys are required) Administrative properties according to Table 13.4</i>
<i>rectangle properties</i>	<i>(Some keys are required) Rectangle properties according to Table 13.5</i>
<i>appearance properties</i>	<i>(Some keys are required) Appearance properties for all Block types according to Table 13.6 and text appearance properties according to Table 13.7 for Textline and Textflow Blocks</i>
<i>text preparation properties</i>	<i>(Optional) Text preparation properties for Textline and Textflow Blocks according to Table 13.8</i>
<i>text formatting properties</i>	<i>(Optional) Text formatting properties for Textline and Textflow Blocks according to Table 13.9</i>
<i>object fitting properties</i>	<i>(Optional) Object fitting properties for Textline, Image, PDF, and Graphics Blocks according to Table 13.10, and fitting properties for Textflow Blocks according to Table 13.11</i>
<i>properties for default contents</i>	<i>(Optional) Properties for default contents according to Table 13.13</i>
<i>Custom</i>	<i>(Dictionary; optional) A dictionary containing key/value pairs for custom properties according to Table 13.14.</i>

A Revision History

<i>Date</i>	<i>Changes</i>
<i>September 21, 2022</i>	▶ <i>Updates for PDFlib 10.0.1</i>
<i>January 25, 2022</i>	▶ <i>Updates for PDFlib 10.0.0</i>
<i>March 29, 2021</i>	▶ <i>Updates for PDFlib 9.3.1</i>
<i>July 14, 2020</i>	▶ <i>Updates for PDFlib 9.3.0</i>
<i>February 01, 2019</i>	▶ <i>Updates for PDFlib 9.2.0</i>
<i>February 01, 2018</i>	▶ <i>Updates for PDFlib 9.1.2</i>
<i>July 24, 2017</i>	▶ <i>Updates for PDFlib 9.1.1</i>
<i>December 15, 2016</i>	▶ <i>Updates for PDFlib 9.1.0</i>
<i>July 27, 2016</i>	▶ <i>Updates for PDFlib 9.0.7</i>
<i>November 23, 2015</i>	▶ <i>Updates for PDFlib 9.0.6</i>
<i>May 18, 2015</i>	▶ <i>Updates for PDFlib 9.0.5</i>
<i>December 16, 2014</i>	▶ <i>Updates for PDFlib 9.0.4</i>
<i>May 14, 2014</i>	▶ <i>Updates for PDFlib 9.0.3</i>
<i>December 17, 2013</i>	▶ <i>Updates for PDFlib 9.0.2</i>
<i>July 24, 2013</i>	▶ <i>Updates for PDFlib 9.0.1</i>
<i>March 12, 2013</i>	▶ <i>Updates for PDFlib 9.0.0</i>
<i>June 09, 2011</i>	▶ <i>Updates for PDFlib 8 VT Edition (internally 8.1.0)</i>
<i>December 09, 2010</i>	▶ <i>Various updates and corrections for PDFlib 8.0.2</i>
<i>September 22, 2010</i>	▶ <i>Various updates and corrections for PDFlib 8.0.1p7</i>
<i>April 13, 2010</i>	▶ <i>Various updates and corrections for PDFlib 8.0.1</i>
<i>December 07, 2009</i>	▶ <i>Updates for PDFlib 8.0.0</i>
<i>April 20, 2010</i>	▶ <i>Minor corrections for PDFlib 7.0.5</i>
<i>March 13, 2009</i>	▶ <i>Various updates and corrections for PDFlib 7.0.4</i>
<i>February 13, 2008</i>	▶ <i>Various updates and corrections for PDFlib 7.0.3</i>
<i>August 08, 2007</i>	▶ <i>Various updates and corrections for PDFlib 7.0.2</i>
<i>February 19, 2007</i>	▶ <i>Various updates and corrections for PDFlib 7.0.1</i>
<i>October 03, 2006</i>	▶ <i>Updates and restructuring for PDFlib 7.0.0</i>
<i>February 15, 2007</i>	▶ <i>Various updates and corrections for PDFlib 6.0.4</i>
<i>February 21, 2006</i>	▶ <i>Various updates and corrections for PDFlib 6.0.3; added Ruby section</i>
<i>August 09, 2005</i>	▶ <i>Various updates and corrections for PDFlib 6.0.2</i>
<i>November 17, 2004</i>	▶ <i>Minor updates and corrections for PDFlib 6.0.1</i> ▶ <i>introduced new format for language-specific function prototypes in chapter 8</i> ▶ <i>added hypertext examples in chapter 3</i>

Date	Changes
June 18, 2004	▶ Major changes for PDFlib 6
January 21, 2004	▶ Minor additions and corrections for PDFlib 5.0.3
September 15, 2003	▶ Minor additions and corrections for PDFlib 5.0.2; added block specification
May 26, 2003	▶ Minor updates and corrections for PDFlib 5.0.1
March 26, 2003	▶ Major changes and rewrite for PDFlib 5.0.0
June 14, 2002	▶ Minor changes for PDFlib 4.0.3 and extensions for the .NET binding
January 26, 2002	▶ Minor changes for PDFlib 4.0.2 and extensions for the IBM eServer edition
May 17, 2001	▶ Minor changes for PDFlib 4.0.1
April 1, 2001	▶ Documents PDI and other features of PDFlib 4.0.0
February 5, 2001	▶ Documents the template and CMYK features in PDFlib 3.5.0
December 22, 2000	▶ ColdFusion documentation and additions for PDFlib 3.0.3; separate COM edition of the manual
August 8, 2000	▶ Delphi documentation and minor additions for PDFlib 3.0.2
July 1, 2000	▶ Additions and clarifications for PDFlib 3.0.1
Feb. 20, 2000	▶ Changes for PDFlib 3.0
Aug. 2, 1999	▶ Minor changes and additions for PDFlib 2.0.1
June 29, 1999	▶ Separate sections for the individual language bindings ▶ Extensions for PDFlib 2.0
Feb. 1, 1999	▶ Minor changes for PDFlib 1.0 (not publicly released)
Aug. 10, 1998	▶ Extensions for PDFlib 0.7 (only for a single customer)
July 8, 1998	▶ First attempt at describing PDFlib scripting support in PDFlib 0.6
Feb. 25, 1998	▶ Slightly expanded the manual to cover PDFlib 0.5
Sept. 22, 1997	▶ First public release of PDFlib 0.4 and this manual

Index

A

- Acrobat Plugin for creating Blocks* 333
- advanced linebreaking* 221
- alpha channel* 165
 - from separate image* 166
 - internal* 166
- ArtBox* 61
- ascender* 136
- asciifile option* 58
- associated files* 69
- attachment password* 66

B

- backslash substitution* 105
- Basic Multilingual Plane* 95
- bindings* 27
- black point compensation* 78
- BleedBox* 61
- blend modes* 88, 91
- blending color space* 90
- Block nesting* 356
- Blocks* 333
 - create with POCA* 371
 - plugin* 333
 - properties* 335
- BMP* 95, 163
- bookmarks with structure* 284
- builtin CJK Fonts* 124
- Byte Order Mark (BOM)* 96, 100
- bytes: see hypertextformat*
- byteserving* 57

C

- C binding* 27
- C++ binding* 30
- capheight* 136
- categories of resources* 51
- CCSID* 102
- change the color of objects* 90
- character metrics* 136
- character references* 105, 106
- characters and glyphs* 95
- Chinese* 103, 150
- chroma key masking* 167
- CIE L*a*b* color space* 78
- CJK (Chinese, Japanese, Korean)*
 - configuration* 104
 - custom fonts* 150
 - Windows code pages* 104

- CJK fonts (builtin)* 124
 - clip* 62
 - clone page boxes* 195
- CMaps* 103
- CMYK color* 73
- color blends* 85
- color fonts* 110
- color profiles* 185
- color spaces* 73
- colorize images* 168
- colorize objects* 91
- COLR color fonts* 110
- commercial license* 13
- content strings* 98
- coordinate system* 58
 - metric* 59
 - top-down* 60
- core fonts* 124
- CropBox* 61
- current point* 62
- currentx and currenty options* 136
- custom encoding* 103
- custom structure types* 271

D

- decode array for image colors* 169
- decode image option* 169
- decolorize objects* 91
- default color spaces* 74, 76
- default coordinate system* 58
- defaultgray/rgb/cmyk color space* 76
- descender* 136
- DeviceN colors* 82
- document part hierarchy* 324, 327
- Document Part Metadata (DPM)* 324, 328
- downsampling* 159
- dpi calculations* 159

E

- EBCDIC* 58
- ebcdicutf8: see hypertextformat*
- editable Watermark* 205
- embedding fonts* 129
- emoji modification sequences* 138
- encapsulated XObjects for PDF/VT* 329
- encapsulation hints for PDF/VT* 325
- encoding*
 - custom* 103
 - fetching from the system* 102

- encryption 66
- environment variable *PDFLIBRESOURCEFILE* 54
- error handling 47
- errorpolicy option 183
- escape sequences 105
- EUDC (end-user defined characters) 109, 151
- evaluation version 11
- exceptions 47
- EXIF JPEG images 161

F

- features of PDFlib 21, 23
- file search 52
- fill 62
- fitbox for Textflow 206
- font name alias 122
- fonts
 - embedding 129
 - metrics 136
 - OpenType 109
 - resource configuration 51
 - selecting symbolic glyph 119
 - SING 112
 - standard latin 124
 - style names for Windows 126
 - subsetting 130
 - TrueType 109
 - TrueType Collection 109
 - Type 3 (user-defined) fonts 113
 - WOFF and WOFF2 112
- form fields: converting to PDFlib Blocks 346
- Form XObjects 64

G

- gaiji characters 113
- get_buffer() 56
- glyph 95
 - availability in font 132
 - glyph id (GID) addressing 114
 - replacement 157
 - selecting from symbol font 119
- glyphlets 112
- gradients 85
- graphics 170

H

- HKS colors 80
- horizontal writing mode 150
- host encoding 101
- host fonts 126
- HSL color representation, hue, saturation, color, luminosity 88
- HTML character references 105
- hypertext strings 99
- hypertextformat option 99

I

- IBM System i 58
- IBM Z 58
- ICC color profiles 185
- iccprofilegray/rgb/cmyk options 76
- Ideographic Variation Sequences (IVS) 154
- image
 - file formats 161
 - inline 114
 - re-using data 159
 - scaling 159
 - transparency 165
- inch 58
- in-core PDF generation 55
- inline image 114, 160
- invert color of objects 91
- invisible text 361
- ISO 14289-1 (PDF/UA-1) 293
- ISO 15930 (PDF/X) 316
- ISO 16612-2 (PDF/VT) 324
- ISO 19005 (PDF/A) 305
- ISO 32000-1 (PDF 1.7) 301
- ISO 32000-2 (PDF 2.0) 303

J

- Japanese 103, 150
- Java binding 31
- Javadoc 33
- JBIG2 162
- JFIF 161
- JPEG 161
 - images in EXIF format 161
- JPEG 2000 161

K

- kerning 136
- Korean 104, 150

L

- Lab color space 78
- language bindings: see bindings
- layers and PDI 185
- leading 136
- line spacing 136
- logging 49
- logical reading order 276
- luminosity mask 91

M

- masked 167
- master password 66
- MediaBox 61
- memory, generating PDF documents in 55
- metric coordinates 59
- metrics 136

millimeters 58
multi-page image files 160

N

name strings 99
NChannel color spaces 83
n-colorant ICC profiles 317
nested Blocks 356
nesting exceptions 28
.NET binding 33
noaccessible 68
noannots 68
noassemble 68
nocopy 68
noforms 68
nohiresprint 68
nomodify 68
noprint 67

O

Objective-C binding 36
OBJR structure element for interactive elements
282, 284
OpenType
collection 109
color fonts 110
font variations 109
fonts 109
layout features 139
orphan lines 216
OTC (OpenType Collection) 150
outline text 361
output intent 318
for PDF/A 308
for PDF/X 319, 320
overline option 138
overprint control 93
owner password 66

P

page 160
page descriptions 58
page formats 61
page size limitations in Acrobat 61
Pantone colors 80
passwords 66
path 62
path objects 62
patterns (tiling) 87
PDF 1.7 301
PDF 1.7 and PDF 2.0 tagsets 261
PDF 2.0 303
PDF import library (PDI) 180
PDF_EXIT_TRY() 29
PDF_get_buffer() 56
PDF/A 305

PDF/UA 293
PDF/VT 324
PDF/X 316
PDFlib Blocks 333
PDFlib features 21, 23
PDFlib Personalization Server (PPS) 333
pdflib.upr 54
PDFLIBRESOURCEFILE environment variable 54
PDI (PDF Import) 180
pdusebox 185
Perl binding 37
permissions 66, 67
permissions password 66
Photoshop CMYK images 164
PHP binding 38
plainmetadata 68
plugin for creating Blocks 333
PNG 161
POCA (PDF Object Creation API)
for creating Blocks 371
for Document Part Metadata (DPM) 328
PPS (PDFlib Personalization Server) 333
print stream order 276
Private Use Area (PUA) 95, 118
Python binding 40

R

raw image data 163
reading order 276
record level for PDF/VT 324
rendering intents 93
resource category 51
resourcefile option 55
RGB color 73
role map for custom structure types 271
rotating objects 59
RPG binding 41
Ruby binding 43

S

sbix color fonts 110
scalable vector graphics (SVG) 170
scaling images 159
scope hints for PDF/VT 328
script-specific linebreaking 221
SearchPath option 52
shadings 85
SING fonts 112, 151
soft mask 91
spot color (separation color space) 79
sRGB color space 75
standard 14 fonts 124
standard structure types 261
standardized variation sequences 153
stencil mask 165, 168
strikeout option 138
stroke 62

strongly structured documents 297
structured bookmarks 284
style names for Windows 126
subpath 62
subscript 136
subsetting 130
superscript 136
SVG 170
 color 176
SVG color fonts 110
symbol font, selecting glyphs 119
system encoding support 102

T

tagset 261
tagset namespaces 271
tagsets 261
templates 64
text metrics 136
text position 136
text style 136
Textflow 205
 fitbox 206
textformat option 99
textrendering option 138
textx and texty options 136
TIFF 162
tiling patterns 87
top-down coordinates 60
transparency
 detect in imported PDF pages 330
 in images 165
 in PDF/VT 329
TrimBox 61
TrueType fonts 109
TTC (TrueType Collection) 109, 150
Type 3 (user-defined) fonts 113

U

underline option 138
Unicode variation sequence (UVS) 153
unique identification of XObjects for PDF/VT 325
units 58
UPR (Unix PostScript Resource) 51
usehypertextencoding option 99
user password 66
user space 58
usercoordinates option 59
user-defined (Type 3) fonts 113
UTF formats 96
utf8/utf16/utf16be/utf16le: see hypertextformat

V

variable fonts 109
variation selectors and variation sequences 153
vector graphics 170
vertical writing mode 150

W

watermark (editable) 205
weakly structured documents 297
widow lines 216
WOFF and WOFF2 fonts 112
writing modes 150

X

xCLR ICC rofiles 317
xheight 136
XMP metadata 68
XObjects 64

Z

ZUGFeRD standard for electronic invoices 306

PDFlib GmbH

Franziska-Bilek-Weg 9
80339 München, Germany
www.pdflib.com
phone +49 • 89 • 452 33 84-0

Licensing contact

sales@pdflib.com

Support

support@pdflib.com *(please include your license number)*

