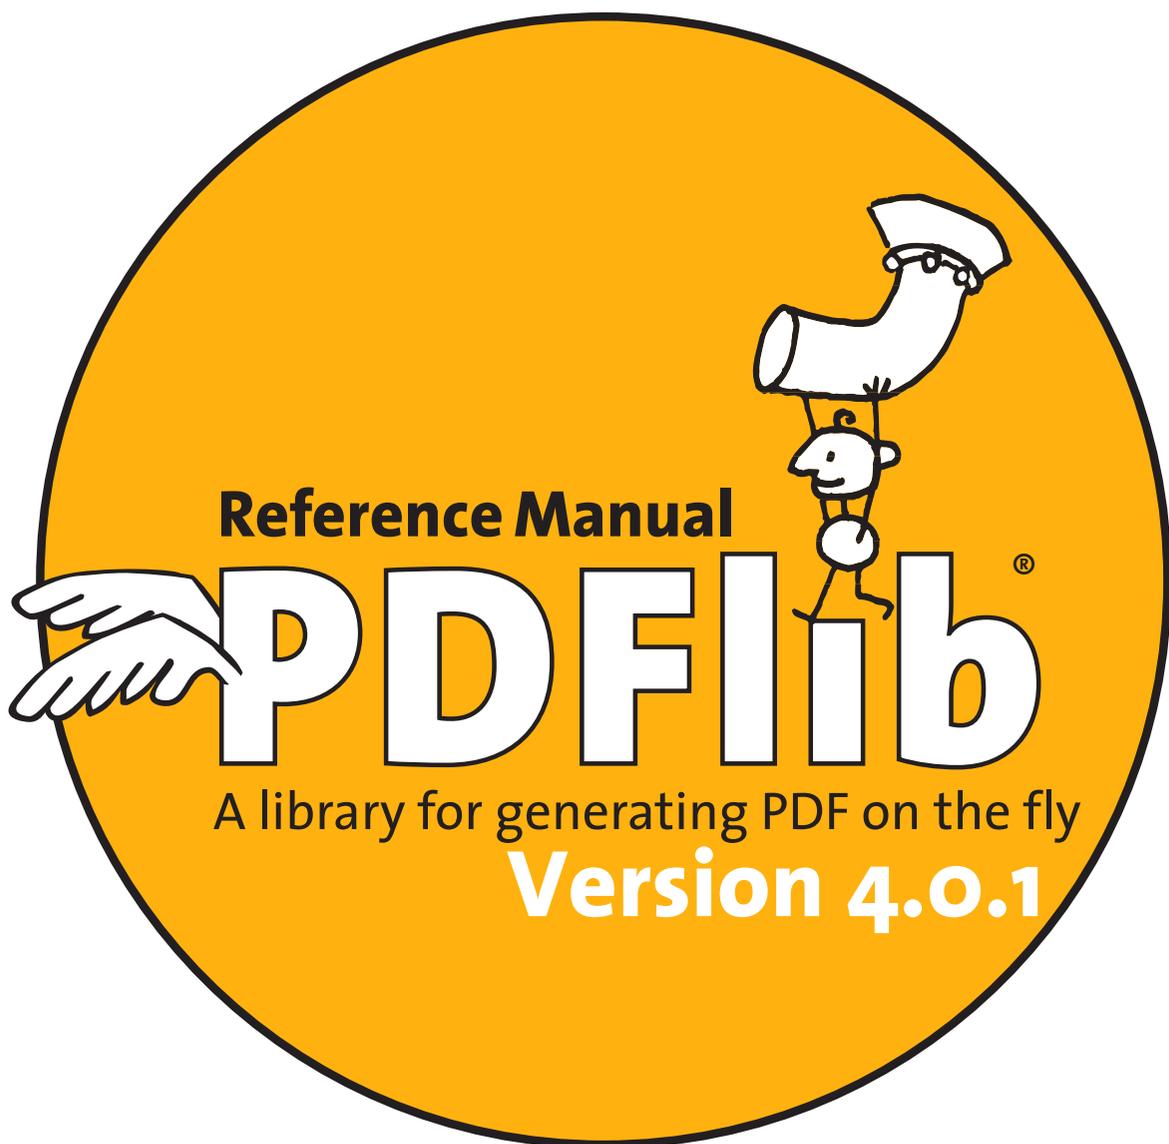


PDFlib GmbH München, Germany



ActiveX/COM edition

www.pdfliib.com

Copyright © 1997–2001 PDFlib GmbH and Thomas Merz. All rights reserved.

PDFlib GmbH
Tal 40, 80331 München, Germany
<http://www.pdflib.com>

phone +49 • 89 • 29 16 46 87
fax +49 • 89 • 29 16 46 86

If you have questions check the PDFlib mailing list and archive at <http://groups.yahoo.com/group/pdflib>

Licensing contact: sales@pdflib.com
Support for commercial PDFlib licensees: 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 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.

Adobe, Acrobat, and PostScript are trademarks of Adobe Systems Inc. AIX, IBM, OS/390, WebSphere, iSeries, and zSeries are trademarks or registered trademarks of International Business Machines Corporation. ActiveX, Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation. Apple, Macintosh and TrueType are trademarks of Apple Computer, Inc. Unicode and the Unicode logo are trademarks of Unicode, Inc. Unix is a trademark of The Open Group. Java and Solaris are a trademark of Sun Microsystems, Inc. Other company product and service names may be trademarks or service marks of others.

*PDFlib contains modified parts of the following third-party software:
PNG image reference library (libpng), Copyright © 1998, 1999, 2000 Glenn Randers-Pehrson
Zlib compression library, Copyright © 1995-1998 Jean-loup Gailly and Mark Adler
TIFFlib image library, Copyright © 1988-1997 Sam Leffler, Copyright © 1991-1997 Silicon Graphics, Inc.*



Author: Thomas Merz
Design and illustrations: Alessio Leonardi
Quality control (manual): Katja Karsunke, Petra Porst, York Karsunke, Rainer Schaaf
Quality control (software): a cast of thousands

Contents

1 Introduction 7

- 1.1 PDFlib Programming 7
- 1.2 PDFlib Features 9
- 1.3 PDFlib Output and Compatibility 10
- 1.4 What's new in PDFlib 4.0? 11

2 PDFlib Language Bindings 12

- 2.1 Overview of the PDFlib Language Bindings 12
 - 2.1.1 What's all the Fuss about Language Bindings? 12
 - 2.1.2 Availability and Platforms 12
 - 2.1.3 The »Hello world« Example 12
 - 2.1.4 Error Handling 12
 - 2.1.5 Version Control 13
 - 2.1.6 Unicode Support 13
 - 2.1.7 Summary of Language Bindings 14
- 2.2 ActiveX/COM Binding 14
 - 2.2.1 How does the ActiveX/COM Binding work? 14
 - 2.2.2 Installing the PDFlib ActiveX Edition 15
 - 2.2.3 Error Handling in ActiveX 17
 - 2.2.4 Version Control in ActiveX 18
 - 2.2.5 Unicode Support in ActiveX 18
 - 2.2.6 Using PDFlib with Active Server Pages 18
 - 2.2.7 Using PDFlib with Visual Basic 22
 - 2.2.8 Using PDFlib with Windows Script Host 24
 - 2.2.9 Using PDFlib with ColdFusion 25
 - 2.2.10 Using PDFlib with Borland Delphi 26
- 2.3 C Binding 28
- 2.4 C++ Binding 28
- 2.5 Java Binding 28
- 2.6 Perl Binding 28
- 2.7 PHP Binding 28
- 2.8 Python Binding 28
- 2.9 RPG Binding 28
- 2.10 Tcl Binding 28

3 PDFlib Programming Concepts 29

- 3.1 General Programming Issues 29
 - 3.1.1 The PDFlib Demo Stamp and Serial Numbers 29

- 3.1.2 PDFlib Program Structure **29**
- 3.1.3 Generating PDF Documents directly in Memory **29**
- 3.1.4 Error Handling **30**
- 3.2 Page Descriptions 31**
 - 3.2.1 Coordinate Systems **31**
 - 3.2.2 Page and Coordinate Limits **33**
 - 3.2.3 Paths and Color **34**
 - 3.2.4 Templates **35**
- 3.3 Text Handling 36**
 - 3.3.1 The PDF Core Fonts **36**
 - 3.3.2 8-Bit Encodings built into PDFlib **37**
 - 3.3.3 Custom Encoding and Code Page Files for 8-Bit Encodings **38**
 - 3.3.4 Hypertext Encoding **41**
 - 3.3.5 PostScript and TrueType Fonts **41**
 - 3.3.6 Resource Configuration and the UPR Resource File **45**
 - 3.3.7 CID Font Support for Japanese, Chinese, and Korean Text **48**
 - 3.3.8 Unicode Support **52**
 - 3.3.9 Text Metrics, Text Variations, and Text Box Formatting **54**
- 3.4 Image Handling 58**
 - 3.4.1 Supported Image File Formats **58**
 - 3.4.2 Code Fragments for Common Image Tasks **59**
 - 3.4.3 Re-using Image Data **61**
 - 3.4.4 Memory Images and External Image References **62**
 - 3.4.5 Image Masks and Transparency **62**
 - 3.4.6 Multi-Page Image Files **64**
- 3.5 PDF Import with PDI 65**
 - 3.5.1 PDI Features and Applications **65**
 - 3.5.2 Using PDI Functions with PDFlib **65**
 - 3.5.3 Acceptable PDF Documents **67**
- 4 PDFlib API Reference 69**
 - 4.1 Data Types, Naming Conventions, and Scope 69**
 - 4.2 General Functions 70**
 - 4.2.1 Setup **70**
 - 4.2.2 Document and Page **71**
 - 4.2.3 Parameter Handling **73**
 - 4.3 Text Functions 74**
 - 4.3.1 Font Handling **74**
 - 4.3.2 Text Output **75**
 - 4.4 Graphics Functions 79**
 - 4.4.1 General Graphics State **79**
 - 4.4.2 Special Graphics State **81**
 - 4.4.3 Path Construction **84**
 - 4.4.4 Path Painting and Clipping **86**

- 4.5 Color Functions 88**
- 4.6 Image Functions 90**
- 4.7 PDF Import Functions 95**
 - 4.7.1 Document and Page 95
 - 4.7.2 Parameter Handling 96
- 4.8 Hypertext Functions 98**
 - 4.8.1 Document Open Action and Open Mode 98
 - 4.8.2 Bookmarks 98
 - 4.8.3 Document Information Fields 99
 - 4.8.4 Page Transitions 99
 - 4.8.5 File Attachments 100
 - 4.8.6 Note Annotations 101
 - 4.8.7 Links 101
 - 4.8.8 Thumbnails 104
- 4.9 Page Size Formats 104**

5 The PDFlib License 105

- 5.1 The »Aladdin Free Public License« 105
- 5.2 The Commercial PDFlib License 105

6 References 106

A Shared Libraries and DLLs 107

B PDFlib Quick Reference 108

C Revision History 112

Index 113



1 Introduction

1.1 PDFlib Programming

What is PDFlib? PDFlib is a library which allows you to generate files in Adobe's Portable Document Format (PDF). PDFlib acts as a backend to your own programs. While you (the programmer) are responsible for retrieving or maintaining the data to be processed, PDFlib takes over the task of generating the PDF code which graphically represents your data. While you must still format and arrange your text and graphical objects, PDFlib frees you from the internals and intricacies of PDF. PDFlib offers many useful functions for creating text, graphics, images and hypertext elements in PDF files.

How can I use PDFlib? PDFlib is available on a variety of platforms, including Unix, Windows, Mac OS, and EBCDIC-based systems such as IBM eServer iSeries 400 and zSeries S/390. PDFlib itself is written in the C language, but it can be also accessed from several other languages and programming environments which are called language

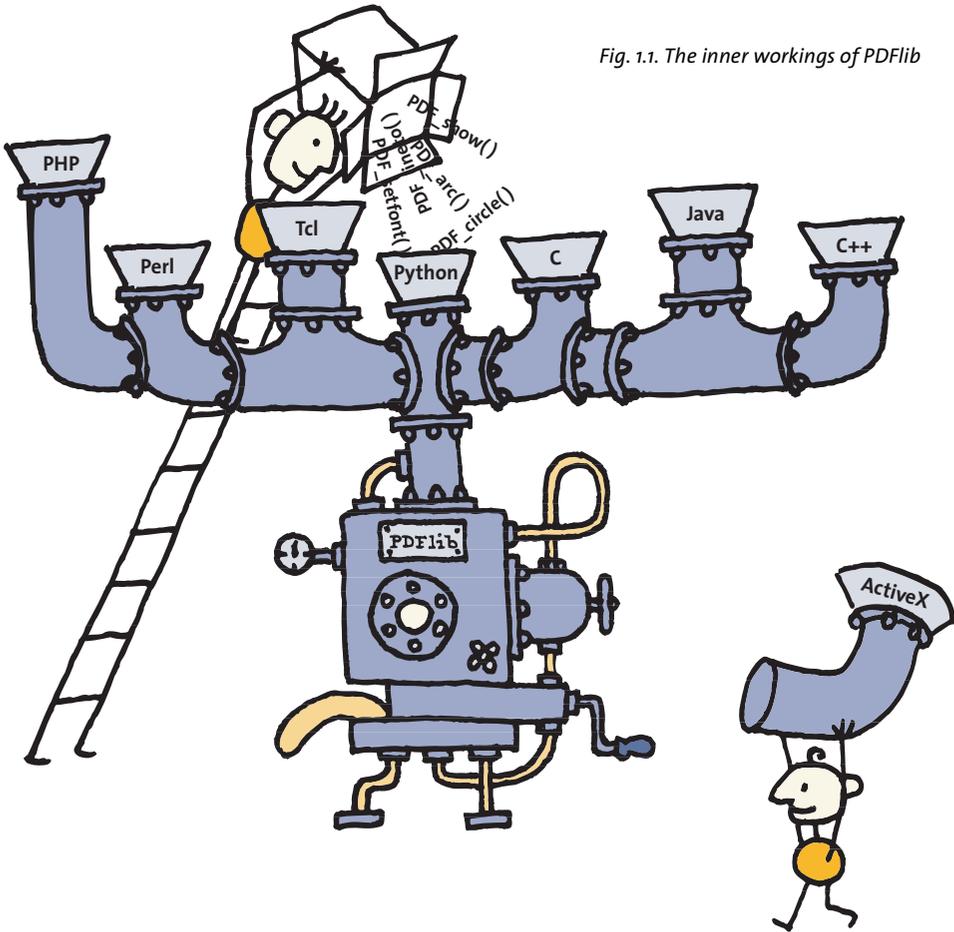


Fig. 1.1. The inner workings of PDFlib

bindings. These language bindings cover all major Web application languages currently in use. The Application Programming Interface (API) is easy to learn, and is identical for all bindings. Currently the following bindings are supported:

- ▶ ActiveX/COM, providing access from Visual Basic, Active Server Pages with VBScript or JScript, Allaire ColdFusion, Borland Delphi, Windows Script Host, and many other environments
- ▶ ANSI C
- ▶ ANSI C++
- ▶ Java, including servlets
- ▶ PHP hypertext processor
- ▶ Perl
- ▶ Python
- ▶ Tcl

What can I use PDFlib for? PDFlib's primary target is creating dynamic PDF within your own software, on the World Wide Web. Similar to HTML pages dynamically generated on the Web server, you can use a PDFlib program for dynamically generating PDF reflecting user input or some other dynamic data, e.g. data retrieved from the Web server's database. The PDFlib approach offers several advantages:

- ▶ PDFlib can be integrated directly in the application generating the data, eliminating the convoluted creation path application–PostScript–Acrobat Distiller–PDF.
- ▶ As an implication of this straightforward process, PDFlib is the fastest PDF-generating method, making it perfectly suited for the Web.
- ▶ PDFlib's thread-safety as well as its robust memory and error handling support the implementation of high-performance server applications.
- ▶ PDFlib is available for a variety of operating systems and development environments.

However, PDFlib is not restricted to dynamic PDF on the Web. Equally important are all kinds of converters from X to PDF, where X represents any text or graphics file format. Again, this replaces the sequence X–PostScript–PDF with simply X–PDF, which offers many advantages for some common graphics file formats like TIFF, GIF, PNG or JPEG. Using such a PDF converter, batch converting lots of text or graphics files is much easier than using the Adobe Acrobat suite of programs.

Requirements for using PDFlib. PDFlib makes PDF generation possible without wading through the 600+ page PDF specification. While PDFlib tries to hide technical PDF details from the user, a general understanding of PDF is useful. In order to make the best use of PDFlib, application programmers should ideally be familiar with the basic graphics model of PostScript (and therefore PDF). However, a reasonably experienced application programmer who has dealt with any graphics API for screen display or printing shouldn't have much trouble adapting to the PDFlib API as described in this manual.

About this manual. This manual describes the API implemented in PDFlib. The function interfaces described in this manual are believed to remain unchanged during future PDFlib development. This manual does not attempt to explain Acrobat features. Please refer to the Acrobat product literature, and the material cited at the end of this manual for further reference.

1.2 PDFlib Features

Table 1.1 lists the major PDFlib API features for generating PDF documents.

Table 1.1. PDFlib features for generating PDF

topic	features
PDF input	<ul style="list-style-type: none">▶ existing PDF documents can be imported with the optional PDF import library (PDI)
PDF output	<ul style="list-style-type: none">▶ PDF documents of arbitrary length, directly in memory (for Web servers) or on disk file▶ arbitrary page size—each page may have a different size▶ compression for text, vector graphics, image data, and file attachments▶ compatibility modes for PDF 1.2, 1.3, and 1.4 (Acrobat 3, 4, and 5)
Vector graphics	<ul style="list-style-type: none">▶ common vector graphics primitives: lines, curves, arcs, rectangles, etc.▶ vector paths for stroking, filling, and clipping▶ grayscale, RGB, CMYK, and spot color for stroking and filling objects▶ pattern fills and strokes▶ efficiently re-use text or vector graphics with templates
Fonts	<ul style="list-style-type: none">▶ text output in different fonts; underlined, overlined, and strikethrough text▶ text column formatting▶ built-in font metrics for PDF's 14 base fonts▶ TrueType and PostScript Type 1 (PFB and PFA file formats) font support with or without font embedding; TrueType fonts can be pulled from the Windows host system▶ support for AFM and PFM PostScript font metrics files▶ library clients can retrieve character metrics for exact formatting▶ on IBM eServer iSeries 400 and zSeries S/390: fetch encodings from the system
Hypertext	<ul style="list-style-type: none">▶ page transition effects such as shades and mosaic▶ nested bookmarks▶ PDF links, launch links (other document types), and Web links▶ document information: four standard fields (Title, Subject, Author, Keywords) plus unlimited number of user-defined info fields (e.g., part number)▶ file attachments and note annotations
Internationalization	<ul style="list-style-type: none">▶ Unicode support (see below)▶ support for a variety of encodings (both built-in and user-defined)▶ CID font and CMap support for Chinese, Japanese, and Korean text▶ support for the Euro character▶ support for international standards and vendor-specific code pages
Unicode	<ul style="list-style-type: none">▶ Unicode support for hypertext features: bookmarks, contents and title of text annotations, document information fields, attachment description, and author name▶ Unicode code pages for TrueType and PostScript fonts▶ Unicode encoding for Japanese, Chinese, and Korean text
Images	<ul style="list-style-type: none">▶ embed GIF (non-interlaced), PNG, TIFF, JPEG, or CCITT raster images▶ Images constructed by the client directly in memory▶ efficiently re-use image data, e.g., for repeated logos on each page▶ transparent (masked) images
Programming	<ul style="list-style-type: none">▶ language bindings for ActiveX/COM, C, C++, Java (including servlets), Perl, PHP, Python, Tcl▶ transparent Unicode handling for ActiveX, Java, and Tcl▶ thread-safe for deployment in multi-threaded server applications▶ configurable error handler and memory management for C and C++▶ exception handling integrated with the host language's native exception handling▶ available for a wide variety of systems, including ASCII- and EBCDIC-based platforms

1.3 PDFlib Output and Compatibility

PDFlib output. PDFlib generates binary PDF output. The PDF output will be compressed with the Flate (also known as ZIP) compression algorithm. Compression can also be deactivated. Compression applies to potentially large items, such as raster image data and file attachments, as well as text and vector operators on page descriptions. The compression speed/output size trade-off can be controlled with a PDFlib parameter.

Acrobat 4 features. Generally, we strive to produce PDF documents which may be used with a wide variety of PDF consumers. PDFlib generates output compatible with Acrobat 3 and higher.

However, certain features either require Acrobat 4, or don't work in Acrobat Reader but only the full Acrobat product. Table 1.2 lists those features. More details can be found at the respective function descriptions.

Table 1.2. PDFlib features which require Acrobat 4

topic	remarks
hypertext	<ul style="list-style-type: none">▶ file attachments are not recognized in Acrobat 3 (require full Acrobat 4)▶ different icons for notes are not recognized in Acrobat 3
page size	▶ Acrobat 4 extends the limits for acceptable PDF page sizes
Unicode	▶ Unicode hypertext doesn't work in Acrobat 3
font	<ul style="list-style-type: none">▶ the Euro symbol is not supported in Acrobat 3▶ Unicode support for TrueType fonts doesn't work in Acrobat 3▶ CID fonts for Chinese, Japanese, and Korean require Acrobat 3J or Acrobat 4
color	▶ the pattern color space is not supported in Acrobat 3 compatibility mode (although patterns can be printed with Acrobat 3, they do not display on screen).
transparency	▶ transparency information is ignored in Acrobat 3
JPEG images	▶ Acrobat 3 supports only baseline JPEG images, but not the progressive flavor
external images	▶ Acrobat 4 (but not the free Acrobat Reader) support external image references via URL. Acrobat 3 is unable to display such referenced images.

Acrobat 3 compatibility mode. Basically, if you don't use the above-mentioned Acrobat 4 features, the generated PDF files will be compatible to Acrobat 3 and 4. However, due to a very subtle compatibility issue with certain output devices, PDFlib also offers a strict Acrobat 3 compatibility mode. In order to understand the problem, we must distinguish between the actual Acrobat viewer version required by a certain PDF file, and the very first line in the file which may read `%PDF-1.2` or `%PDF-1.3` for Acrobat 3 and Acrobat 4-generated files, respectively. It's important to know that Acrobat 3 viewers open files starting with the `%PDF-1.3` line without any problem, provided the file doesn't use any Acrobat 4 feature. This is the basis of PDFlib's multi-version compatibility approach.

However, some PDF consumers other than Acrobat implement a much stricter way of version control: they simply reject all files starting with the `%PDF-1.3` line, regardless of whether the actual content requires a PDF 1.2 or PDF 1.3 interpreter. For example, some Efi RIPs for high-speed digital printing machines are known to (mis-)behave in this manner. In order to work around this problem, PDFlib offers a strict Acrobat 3 compatibility mode in which a `%PDF-1.2` header is emitted, and Acrobat 4 features are disabled.

Note again that it is not necessary to use PDFlib's strict Acrobat 3 compatibility mode only to make sure the PDF files can be read with Acrobat 3 – this will automatically be the case if you refrain from using the above-mentioned Acrobat 4 features. The strict mode is only required for those rare situations where you have to deal with one of those broken PDF-enabled RIPs.

Acrobat 5 compatibility. PDFlib accepts Acrobat 5 PDF files for import, and will generate Acrobat 5 features in the future. Output compatibility may be set to PDF 1.4 (=Acrobat 5) if Acrobat 5 PDF files are to be imported into the generated document. The PDF import library PDI fully supports PDF 1.4 (see Section 3.5.2, »Using PDI Functions with PDFlib«).

1.4 What's new in PDFlib 4.0?

The following list gives a quick overview of new features in PDFlib 4.0 + PDI:

- ▶ support for TrueType fonts on all platforms
- ▶ TrueType host font support on Windows (grab fonts directly from the system)
- ▶ Unicode-based code pages
- ▶ CMYK and spot color
- ▶ template feature
- ▶ pattern tiling
- ▶ fast pass-through mode for PNG images
- ▶ PHP language binding
- ▶ completely integrates all auxiliary libraries (zlib, libpng, libtiff)
- ▶ separate PDF import library (PDI) for dealing with existing PDF documents
- ▶ on IBM eServer iSeries 400 and zSeries S/390: fetch encodings from the system
- ▶ additional language binding for ILE-RPG

2 PDFlib Language Bindings

2.1 Overview of the PDFlib Language Bindings

2.1.1 What's all the Fuss about Language Bindings?

While the C programming language has been one of the cornerstones of systems and applications software development for decades, a whole slew of other languages have been around for quite some time which are either related to new programming paradigms (such as C++), open the door to powerful platform-independent scripting capabilities (such as Perl, Tcl, and Python), promise a new quality in software portability (such as Java), or provide the glue among many different technologies while being platform-specific (such as ActiveX/COM).

Note This edition of the manual discusses only the PDFlib ActiveX editions. All other language bindings are discussed in detail in a separate edition of the manual. Omitted sections of this manual are not relevant for users of the PDFlib ActiveX edition.

2.1.2 Availability and Platforms

All PDFlib features are available on all platforms and in all language bindings (with a few minor exceptions which are noted in the manual). Given the broad range of platforms and languages (let alone different versions of both) supported by PDFlib, it shouldn't be much of a surprise that not all combinations of platforms, languages, and versions thereof can be tested. However, we strive to make PDFlib work with the latest available versions of the respective environments. Table 2.1 lists the language/platform combinations we used for testing.

PDFlib on embedded systems. It shall be noted that PDFlib can also be used on embedded systems, and has been ported to the Windows CE and EPOC environments as well as custom embedded systems. For use with restricted environments certain features are configurable in order to reduce PDFlib's overall memory footprint. If you are interested in details please contact us via sales@pdflib.com.

2.1.3 The »Hello world« Example

Being a well-known programming classic, the »Hello, world!« example will be used for our examples. It uses PDFlib to generate a one-page PDF file with some text on the page. In the following sections, the »Hello, world!« sample will be shown for all supported language bindings. The code for all language samples is contained in the PDFlib distribution. The distribution contains simple examples for text, vector, and image handling as well as PDF import for all supported language bindings.

2.1.4 Error Handling

PDFlib provides a sophisticated means for dealing with different kinds of programming and runtime errors. In order to allow for smooth integration to the respective language environment, PDFlib's error handling is integrated into the language's native way of dealing with exceptions. Basically, C and C++ clients can install custom code which is

Table 2.1. Tested language and platform combinations

<i>language</i>	<i>Unix (Linux, Solaris, HP-UX, AIX a.o.)</i>	<i>Windows</i>	<i>Mac OS (Classic)</i>	<i>IBM eServer iSeries 400, zSeries S/390</i>
<i>ActiveX/COM</i>	–	<i>ASP (PWS, IIS 4 and 5) WSH (VBScript 5, JScript 5) Visual Basic 6.0 Borland Delphi 5 Allaire ColdFusion 4.5</i>	–	–
<i>ANSI C</i>	<i>gcc and other ANSI C compilers</i>	<i>Microsoft Visual C++ 6.0 Metrowerks CodeWarrior 5.3 Borland C++ Builder 5</i>	<i>Metrowerks CodeWarrior 5.3</i>	<i>IBM c89</i>
<i>ANSI C++</i>	<i>gcc and other ANSI C++ compilers</i>	<i>Microsoft Visual C++ 6.0 Metrowerks CodeWarrior 5.3</i>	<i>Metrowerks CodeWarrior 5.3</i>	<i>IBM c89</i>
<i>Java</i>	<i>Sun JDK 1.2.2 and 1.3 IBM JDK 1.1.8 Inprise JBuilder 3.5 Kaffe OpenVM 1.0.5</i>	<i>Sun JDK 1.1.8, 1.2.2, and 1.3 Inprise JBuilder 3.5 and 4 Allaire JRun 3.0 Allaire ColdFusion 4.5</i>	<i>MRJ 2.2, based on JDK 1.1.8</i>	<i>JDK 1.1</i>
<i>Perl</i>	<i>Perl 5.005, 5.6</i>	<i>ActivePerl 5.005 and 5.6</i>	<i>MacPerl 5.2.0r4, based on Perl 5.004</i>	–
<i>PHP</i>	<i>PHP 4.04 and 4.05</i>	<i>PHP 4.04 and 4.05</i>	–	–
<i>Python</i>	<i>Python 1.6 and 2.0</i>	<i>Python 1.6 and 2.0</i>	<i>Python 2.0</i>	–
<i>RPG</i>	–	–	–	<i>V3R7Mo and above</i>
<i>Tcl</i>	<i>Tcl 8.3.2 and 8.4a2</i>	<i>Tcl 8.3.2 and 8.4a2</i>	<i>Tcl 8.3.2</i>	–

called when an error occurs. Other language bindings use the existing exception machinery provided by all modern languages. More details on PDFlib's exception handling can be found in Section 3.1.4, »Error Handling«. The sections on error handling in this chapter cover the language-specific details for the supported environments.

2.1.5 Version Control

Taking into account the rapid development cycles of software in general, and Internet-related software in particular, it is important to allow for future improvements without breaking existing clients. In order to achieve compatibility across multiple versions of the library, PDFlib supports several version control schemes depending on the respective language. If the language supports a native versioning mechanism, PDFlib seamlessly integrates it so the client doesn't have to worry about versioning issues except making use of the language-supplied facilities. In other cases, when the language doesn't support a suitable versioning scheme, PDFlib supplies its own major, minor, and revision version number at the interface level. These may be used by the client in order to decide whether the given PDFlib implementation is acceptable, or should be rejected because a newer version is required.

2.1.6 Unicode Support

PDFlib supports Unicode for a variety of features (see Section 3.3.8, »Unicode Support« for details). The language bindings, however, differ in their native support for Unicode. If a given language binding supports Unicode strings, the respective PDFlib language wrapper is aware of the fact, and automatically deals with Unicode strings in the correct way.

2.1.7 Summary of Language Bindings

For easy reference, Table 2.2 summarizes important features of the PDFlib language bindings. More details can be found in the respective sections of this manual.

Table 2.2. Summary of the language bindings

<i>language</i>	<i>custom error handling</i>	<i>Unicode conversion</i>	<i>version control</i>	<i>thread-safe</i>	<i>EBCDIC-safe</i>
<i>COM/ActiveX</i>	<i>COM exceptions</i>	<i>yes</i>	<i>Class ID and ProgID</i>	<i>yes (both-threading)</i>	<i>–</i>
<i>C</i>	<i>client-supplied error handler</i>	<i>–</i>	<i>manually</i>	<i>yes</i>	<i>yes</i>
<i>C++</i>	<i>client-supplied error handler</i>	<i>–</i>	<i>manually</i>	<i>yes</i>	<i>yes</i>
<i>Java</i>	<i>Java exceptions</i>	<i>yes</i>	<i>automatically</i>	<i>yes</i>	<i>yes</i>
<i>Perl</i>	<i>Perl exceptions</i>	<i>–</i>	<i>via package mechanism</i>	<i>–</i>	<i>–</i>
<i>PHP</i>	<i>PHP warnings</i>	<i>–</i>	<i>manually</i>	<i>yes</i>	<i>–</i>
<i>Python</i>	<i>Python exceptions</i>	<i>–</i>	<i>manually</i>	<i>–</i>	<i>–</i>
<i>RPG</i>	<i>–</i>	<i>–</i>	<i>manually</i>	<i>yes</i>	<i>yes</i>
<i>Tcl</i>	<i>Tcl exceptions</i>	<i>yes (Tcl 8.2 and above)</i>	<i>via package mechanism</i>	<i>yes</i>	<i>–</i>

2.2 ActiveX/COM Binding

2.2.1 How does the ActiveX/COM Binding work?

COM (Component Object Model)¹, developed by Microsoft, is a powerful mechanism for reusing software components regardless of the programming language on the client side (the user of the component) or the server side (the actual implementation of a component). In theory, COM is even a platform-independent binary standard which allows clients to communicate with servers within the same process, on the same machine, or a networked machine. In practice, however, COM is basically a standard for the Windows environment (although attempts have been made to port COM to other platforms).

ActiveX is built on Microsoft's COM technology, and used primarily to develop interactive content for the World Wide Web, although it can be used in desktop and other applications. The reusable software components are called ActiveX controls (formerly known as OLE controls or OCX). Although ActiveX burdens the developer with a variety of specific technologies, terms, and troublesome issues (such as type libraries, registration, an assortment of threading models, historical jettison, to name but a few), ActiveX users are rewarded with tight integration and almost universal usability (if you happen to live in the Windows universe).

Since PDFlib is pure component-ware, the library naturally lends itself to an ActiveX implementation for Windows deployment. The ActiveX implementation of PDFlib is built as a wrapper DLL around the core PDFlib DLL. The wrapper DLL calls the PDFlib core functions and is responsible for communicating with the underlying COM machinery, registration and type library issues, COM exception handling, memory management,

1. See <http://www.microsoft.com/com> for more information about COM and ActiveX

and Unicode string conversions. This parallels other PDFlib bindings where we strive to make a strict distinction between core functionality and the language wrapper. The PDFlib ActiveX wrapper can technically be characterized as follows (don't worry if you are not familiar with all of these terms – they are not required for using PDFlib):

- ▶ PDFlib acts as a Win32 in-process ActiveX server component (also known as an automation server) without any user interface.
- ▶ PDFlib is a »both-threaded« component, i.e., it is treated as both an apartment-threaded as well as a free-threaded component. In addition, PDFlib aggregates a free-threaded marshaler. In simple terms, clients can use the PDFlib object directly (instead of going through a proxy/stub pair) which boosts performance.
- ▶ PDFlib is fully Unicode-aware.
- ▶ The PDFlib binary *pdflib_com.dll* is a self-registering DLL with a type library.
- ▶ PDFlib is stateless, i.e., method parameters are used instead of properties.
- ▶ PDFlib's dual interface supports both early and late binding.
- ▶ PDFlib supports rich error information.

Note PDFlib is currently not MTS-aware (Microsoft Transaction Server).

2.2.2 Installing the PDFlib ActiveX Edition

PDFlib can be deployed in all environments that support ActiveX components. We will demonstrate our examples in several specific environments:

- ▶ Visual Basic
- ▶ Active Server Pages (ASP) with JScript and VBScript
- ▶ Windows Script Host (WSH) with JScript and VBScript
- ▶ Allaire ColdFusion
- ▶ Borland Delphi

Active Server Pages and Windows Script Host both support JScript and VBScript. Since the scripts are nearly identical, however, we do not demonstrate all combinations here. In addition, there are many other ActiveX-aware development environments available – Microsoft Visual C++, Borland C++ Builder, PowerBuilder, to name but a few. PDFlib also works in Visual Basic for Applications (VBA).

Requirements for installing PDFlib. Installing PDFlib is an easy and straight-forward process. Please note the following:

- ▶ PDFlib runs with all Windows versions. No specific operating system version or service pack is required.
- ▶ If you install on an NTFS partition all PDFlib users must have read permission to the installation directory, and execute permission to ...*PDFlib\bin\pdflib_com.dll*.
- ▶ The installer must have write permission to the system registry. Administrator or Power Users group privileges will usually be sufficient.

What the PDFlib ActiveX installer does. The installation program supplied for the PDFlib ActiveX component automatically takes care of all issues related to using PDFlib with ActiveX. For the sake of completeness, the following describes the runtime environment required for using PDFlib (this is taken care of by the installation routine):

- ▶ The PDFlib ActiveX DLL *pdflib_com.dll* is copied to the installation directory.
- ▶ The PDFlib ActiveX DLL must be registered with the Windows registry. The installer uses the self-registering PDFlib DLL to achieve the registration.

- ▶ If a licensed version of PDFlib is installed, the serial number is entered in the system.

While the *Runtime* installation option performs the above steps, the *Full* installation option additionally copies documentation and sample files to the installation directory.

Redistributing the PDFlib ActiveX component. Developers who obtained a redistributable runtime license and wish to redistribute the PDFlib ActiveX component along with their own product must either ship the complete PDFlib installation and run the PDFlib installer as part of their product's setup process, or do all of the following:

- ▶ Integrate the files of the PDFlib *Runtime* installation option in their own installation (see also »Silent install« below). The list of files required by PDFlib can easily be determined by looking at the PDFlib installation directory, since this is the only place where the PDFlib installer places any files. Shipping the supplied AFM files is only required if the core 14 fonts are intended to be used with encodings other than *winansi*, or *builtin* in the case of the Symbol and ZapfDingbats fonts (see Section 3.3.1, »The PDF Core Fonts«).
- ▶ Take care of the necessary PDFlib registry keys. This can be accomplished by completing the entries in the supplied registration file template *pdflib.reg*, and using it during the installation process of your own product.
- ▶ *pdflib_com.dll* must be called for self-registration (e.g., using the *regsvr32* utility). If you use InstallShield for building your installer this can most easily be accomplished by setting the *Self-Registered* property to *true* for the InstallShield group which contains *pdflib_com.dll*.
- ▶ Supply your serial number at runtime using PDFlib's *set_parameter()* function, supplying *serial* as first parameter, and the actual serial string as second parameter (see also Section 3.1.1, »The PDFlib Demo Stamp and Serial Numbers«

```
oPDF.set_parameter("serial", "...your serial string...")
```

Silent install. When PDFlib must be redistributed as part of another software package, or must be deployed on a large number of machines which are administered by tools such as SMS, manually installing PDFlib on each machine may be cumbersome. In such cases PDFlib can also be installed automatically without any user intervention.

The PDFlib installer has been created with InstallShield, and supports InstallShield's silent install feature. A normal (non-silent) installation receives the necessary input from the user in dialog boxes. A silent installation, however, doesn't prompt the user for input. Instead, it gets its input from a special file called the InstallShield silent response (*.iss*) file. A response file is a text file containing information similar to what a user would enter as responses to dialog boxes when running a normal setup. The response file can most easily be prepared by running an interactive installation once, with the installer recording all responses.

Proceed as follows for a silent (non-interactive) installation of PDFlib:

- ▶ Use WinZip to unpack the PDFlib ActiveX installer files into a directory;
- ▶ Run *setup -r* from this directory, and fill out all dialog boxes with the exact entries you wish to use later in the silent installation;
- ▶ Locate the response file *setup.iss* in your Windows directory, and copy it to the directory containing the installer files. You can adjust some of the values in the response file with a text editor if required.
- ▶ To do a silent install run *setup -s* in the directory containing the installer files.

Deploying the PDFlib ActiveX on an ISP's server. Installing software on a server hosted by an Internet Service Provider (ISP) is usually more difficult than on a local machine, since ISPs are often very reluctant when customers wish to install some new software. PDFlib is very ISP-friendly since it doesn't pollute neither the Windows directory nor the registry:

- ▶ Only a single DLL is required, which may live in an arbitrary directory as long as it is properly registered using the *regsvr32* utility.
- ▶ By default only a few private registry entries are created which are located under the *HKEY_LOCAL_MACHINE\SOFTWARE\PDFlib* registry hive. These entries can be manually created if required (see above).
- ▶ If so desired, PDFlib may even be used without any private registry entries. The user must compensate for the entries by using appropriate calls to the *set_parameter()* function for setting the *prefix*, *resourcefile*, and *serial* parameters.

See the section »Redistributing the PDFlib ActiveX component« above for a list of files.

2.2.3 Error Handling in ActiveX

Error handling for the PDFlib Active component is done according to COM conventions: when a PDFlib-internal exception occurs, a COM exception is raised and furnished with the PDFlib error code and a clear-text description of the error. In addition the memory allocated by the PDFlib object is released. Table 2.3 lists all COM errors thrown by PDFlib along with the corresponding PDFlib exceptions. The COM exception may be caught and handled in the PDFlib client in whichever way the client environment supports for handling COM errors.

Table 2.3. COM error codes raised by PDFlib

PDFlib error name	decimal value	hexadecimal value	explanation
<i>MemoryError</i>	-2147220991	&H80040201	<i>not enough memory</i>
<i>IOError</i>	-2147220990	&H80040202	<i>input/output error, e.g. disk full</i>
<i>RuntimeError</i>	-2147220989	&H80040203	<i>wrong order of PDFlib function calls</i>
<i>IndexError</i>	-2147220988	&H80040204	<i>array index error</i>
<i>TypeError</i>	-2147220987	&H80040205	<i>argument type error</i>
<i>DivisionByZero</i>	-2147220986	&H80040206	<i>division by zero</i>
<i>OverflowError</i>	-2147220985	&H80040207	<i>arithmetic overflow</i>
<i>SyntaxError</i>	-2147220984	&H80040208	<i>syntactical error</i>
<i>ValueError</i>	-2147220983	&H80040209	<i>a value supplied as argument to PDFlib is invalid</i>
<i>SystemError</i>	-2147220982	&H8004020A	<i>PDFlib internal error</i>
<i>NonfatalError</i>	-2147220981	&H8004020B	<i>a non-fatal problem was detected</i>
<i>UnknownError</i>	-2147220980	&H8004020C	<i>other error</i>

The error codes used in COM are 32-bit values with the highest bit set, which makes them look like very large negative numbers. PDFlib conforms to the COM conventions, and returns error codes in the range which is reserved for application-defined errors. More specifically, the error codes are constructed as follows:

COM error code = &H80040000 + &H200 + (PDFlib error code)

(The first hexadecimal number is equal to Visual Basic's *vbObjectError* constant, the second is the Microsoft-suggested offset for component-specific errors.)

Table 2.3 lists all PDFlib error names along with the decimal and hexadecimal error codes. A more detailed discussion of PDFlib's exception mechanism can be found in Section 3.1.4, »Error Handling«. Fortunately, ActiveX programmers need not deal with these error numbers directly since the *PDFlib_com* type library provides symbolic constants in the *PDFlib_com.Errors* class.

2.2.4 Version Control in ActiveX

Instead of simple major and minor version numbers, COM implements the concept of a globally unique identifier (GUID) for a class ID which uniquely describes a particular programming interface. Instead of messing around with different version numbers, a new software release may decide whether or not to actually support a certain interface identified via its GUID.

PDFlib_com, being an ActiveX component, makes use of the class ID mechanism. The GUID for *PDFlib_com* is contained in its type library (which in turn is contained in *pdfplib_com.dll*), and in the Windows registry. Since PDFlib is registered under both the generic program identifier (ProgID) *PDFlib_com.PDF*, as well as a version-specific ProgID, users will rarely have to deal with the GUID directly.

2.2.5 Unicode Support in ActiveX

32-bit versions of ActiveX/COM support Unicode natively. The ActiveX language wrapper automatically converts all COM strings to Unicode or ISO Latin 1 (PDFDocEncoding), as appropriate. ActiveX's Unicode-awareness, however, may lead to subtle problems regarding 8-bit encodings (such as *winansi*) and Unicode characters in literal strings. More details on this issue can be found in Section 3.3.8, »Unicode Support«. If you want to use PDFlib's unicode support with ActiveX you must enable unicode mode by setting the *nativeunicode* parameter to *true* (see examples in the next sections).

2.2.6 Using PDFlib with Active Server Pages

Special considerations for Active Server Pages¹. When using external files (such as image files) ASP's *MapPath* facility must be used in order to map path names on the local disk to paths which can be used within ASP scripts. Take a look at the ASP samples supplied with PDFlib, and the ASP documentation if you are not familiar with *MapPath*. Don't use absolute path names in ASP scripts since these may not work without *MapPath*.

In order to access files on another machine from within PDFlib ASP scripts, don't use UNC path names since these won't work in ASP. Instead, add a new virtual directory in IIS. In the administration tool's home directory tab use the option *a share located on another computer*, and supply the UNC name. The remote directory will now be available as a virtual directory in IIS.

The directory containing your ASP scripts must have execute permission, and also write permission unless the in-core method for generating PDF is used (the supplied ASP samples use in-core PDF generation).

1. Active Server Pages is a technology for executing server-side scripts in a variety of languages. It is available with Microsoft Web servers, and several other server products. More information about ASP can be found at <http://msdn.microsoft.com/workshop/server/default.asp>.

You can improve the performance of COM objects such as *PDFlib_com* on Active Server Pages by instantiating the object outside the actual script code on the ASP page, effectively giving the object session scope instead of page scope. More specifically, instead of using *CreateObject* (as shown in the example in the next section)

```
<%@ LANGUAGE = "JavaScript" %>
<%
    var oPDF;
    oPDF = Server.CreateObject("PDFlib_com.PDF");
    if (oPDF.open_file("file.pdf") == -1)
        ...
```

use the *OBJECT* tag with the *RUNAT*, *ID*, and *ProgID* attributes to create the *PDFlib_com* object:

```
<OBJECT RUNAT=Server ID=oPDF ProgID="PDFlib_com.PDF"> </OBJECT>

<%@ LANGUAGE = "JavaScript" %>
<%
    if (oPDF.open_file("file.pdf") == -1)
        ...
```

You can boost performance even more by applying this technique to the *global.asa* file, using the *Scope=Application* attribute, thereby giving the object application scope. Additional ASP examples can be found in the *PDFlib* distribution.

The »Hello world« example for Active Server Pages (ASP) with JScript¹. Unlike the other examples we do not create a PDF output file in the ASP examples. Instead, we generate the PDF data in memory and directly send it to the client via HTTP. This technique is much more appropriate for a Web server environment.

```
<%@ LANGUAGE = "JavaScript" %>
<%
    var font;
    var oPDF;

    oPDF = Server.CreateObject("PDFlib_com.PDF");

    if (oPDF == null) {
        Response.write("Couldn't create PDFlib object!");
        Response.end();
    }

    // Open new PDF file
    oPDF.open_file("");

    oPDF.set_info("Creator", "hello.js.asp");
    oPDF.set_info("Author", "Thomas Merz");
    oPDF.set_info("Title", "Hello, world (ActiveX/ASP/JScript)!");

    // start a new page
    oPDF.begin_page(595, 842);

    font = oPDF.findfont("Helvetica-Bold", "winansi", 0);
```

1. JScript is an extension of ECMAScript (see <http://www.ecma.ch/stand/ecma-262.htm>) which in turn is based on Netscape's JavaScript. For more information on JScript see <http://msdn.microsoft.com/scripting/jscript/default.htm>.

```

oPDF.setfont(font, 24);

oPDF.set_text_pos(50, 700);
oPDF.show("Hello, world!");
oPDF.continue_text("(says ActiveX/ASP/JScript)");

oPDF.end_page();
oPDF.close();

Response.Expires = 0;
Response.Buffer = true;
Response.ContentType = "application/pdf";
Response.Addheader("Content-Disposition", "inline; filename=" +
    "hello.js.asp.pdf");

Response.BinaryWrite(oPDF.get_buffer());
Response.End();
%>

```

The »Hello world« example for Active Server Pages (ASP) with VBScript.

```

<%@ LANGUAGE = VBScript %>
<%
    Option Explicit

    Dim font
    Dim oPDF
    Dim buf

    Set oPDF = Server.CreateObject("PDFlib_com.PDF")

    if not isObject(oPDF) Then
        Response.write "Couldn't create PDFlib object!"
        Response.end
    End If

    ' Open new PDF file
    oPDF.open_file ""

    oPDF.set_info "Creator", "hello.vbs.asp"
    oPDF.set_info "Author", "Thomas Merz"
    oPDF.set_info "Title", "Hello, world (Active X/ASP/VBScript)!"

    ' start a new page
    oPDF.begin_page 595, 842

    font = oPDF.findfont("Helvetica-Bold", "winansi", 0)

    oPDF.setfont font, 24

    oPDF.set_text_pos 50, 700
    oPDF.show "Hello, world!"
    oPDF.continue_text "(says ActiveX/ASP/VBScript)"

    oPDF.end_page
    oPDF.close

```

```

buf = oPDF.get_buffer()

Response.Expires = 0
Response.Buffer = true
Response.ContentType = "application/pdf"
Response.Addheader "Content-Disposition", "inline; filename=" &
    "hello.vbs.asp.pdf"
Response.Addheader "Content-Length", LenB(buf)
Response.BinaryWrite(buf)
Response.End()
Set oPDF = nothing
%>

```

Error handling in JScript. JScript 5.0¹ adds structured exception handling to the language which looks similar to C++ or Java, with the difference that JScript exceptions cannot be typed, and only a single clause can deal with an exception. Detecting an exception and acting upon it is achieved with a *try ... catch* clause:

```

try {
    ...some PDFlib instructions...
} catch (exc) {
    Response.write("Error " + exc.number + ": " + exc.description);
    Response.end();
}

```

Note Due to some problem with JScript's integer handling it's impossible to directly compare exception numbers with hexadecimal values. Comparison with decimal values, however, works fine.

Error handling in VBScript. Unfortunately, VBScript doesn't have any means for catching errors, but only for ignoring them. For this reason one has to periodically check the *ERR* object in order to see whether something went wrong in one of the previous calls to the ActiveX component. VBScript's missing *On Error GoTo* clause has the major drawback that the script code is either cluttered with calls to the error checking routine, or subsequent errors may happen between the first error and the next invocation of the error checking routine:

```

On Error Resume Next
Err.Clear

...some PDFlib instructions...

CheckPDFError

...more PDFlib instructions...

Sub CheckPDFError()
    If Err.number <> 0 then
        WScript.Echo "Error " & Hex(Err.number) & ": " & Err.description
        Err.Clear
    End If
end Sub

```

¹ JScript 5.0 is available with Microsoft Internet Explorer 5.0 and Microsoft Internet Information Services 5.0

Unicode support in JScript. JScript supports Unicode internally. Unicode characters can be written directly into string literals using a Unicode-aware text editor; entered with an escape sequence such as

```
oPDF.set_parameter("nativeunicode", "true");  
Unicodetext = "\u039B\u039F\u0393\u039F\u03A3";
```

or constructed from numerical values using the *String.fromCharCode* method:

```
Unicodetext = String.fromCharCode(0x39B, 0x39F, 0x393, 0x39F, 0x3A3);
```

Unicode support in VBScript. VBScript supports Unicode internally. Similar to Visual Basic, Unicode strings can be constructed from numerical values using the *ChrW* function:

```
oPDF.set_parameter "nativeunicode", "true"  
Unicodetext = ChrW(&H39B) & ChrW(&H39F) & ChrW(&H393) & ChrW(&H39F) & ChrW(&H3A3)
```

2.2.7 Using PDFlib with Visual Basic

Special considerations for Visual Basic¹. When it comes to leveraging external ActiveX components, Visual Basic supports both early (compile-time) and late (run-time) binding. Although both types of bindings are possible with PDFlib, early binding is heavily recommended. It is achieved by performing the following steps:

- ▶ Create a reference from your VB project to PDFlib via *Project, References...*, and select the *PDFlib_com* control.
- ▶ Declare object variables of type *PDFlib_com.PDF* instead of the generic type *Object*:

```
Dim oPDF As PDFlib_com.PDF  
Set oPDF = CreateObject("PDFlib_com.PDF") ' or: Set oPDF = New PDFlib_com.PDF
```

Creating a reference and using early binding has several advantages:

- ▶ VB can check the code for spelling errors.
- ▶ IntelliSense (automatic statement completion) and context-sensitive help are available.
- ▶ The VB object browser shows all PDFlib methods along with their parameters and a short description.
- ▶ VB programs run much faster with early binding than with late binding.

PDFlib programming in Visual Basic is straightforward, with one exception. Due to a Microsoft-confirmed bug (pardon: an »issue«) in Visual Basic 6 several PDFlib functions cannot be used directly since VB erroneously overrides PDFlib method names with some built-in methods of VB. For example, the following cannot be successfully compiled in VB 6:

```
oPDF.circle 10, 10, 30
```

In order to work around this problem, Microsoft technical support came up with the following suggestion:

```
oPDF.[circle] 10, 10, 30
```

¹ *Visual Basic is a commercial product of Microsoft. For more information see <http://msdn.microsoft.com/vbasic/prodinfo>.*

Putting the critical method name in brackets seems to do the trick. From all PDFlib functions only the following are affected by this problem:

```
circle  
scale
```

The data type *integer*, as used in the PDFlib ActiveX component, is a signed 32-bit quantity. In Visual Basic this corresponds to the *long* data type. Therefore, when the PDFlib API reference calls for an *int* type argument, Visual Basic programmers should translate this to *long* (although VB will correctly translate if *int* values are supplied).

The »Hello world« example in Visual Basic.

```
Attribute VB_Name = "hello"  
Option Explicit  
  
Sub main()  
    Dim ret As Long, font As Long  
    Dim oPDF As PDFlib_com.PDF  
  
    Set oPDF = New PDFlib_com.PDF  
  
    ' Open new PDF file  
    ret = oPDF.open_file("hello_ax_vb.pdf")  
    If (ret = -1) Then  
        MsgBox "Couldn't open PDF file!"  
    End If  
End If  
  
oPDF.set_info "Creator", "hello.bas"  
oPDF.set_info "Author", "Thomas Merz"  
oPDF.set_info "Title", "Hello, world (ActiveX/VB)!"  
  
' start a new page  
oPDF.begin_page 595, 842  
  
font = oPDF.findfont("Helvetica-Bold", "winansi", 0)  
  
oPDF.setfont font, 24  
  
oPDF.set_text_pos 50, 700  
oPDF.show "Hello, world!"  
oPDF.continue_text "(says ActiveX/VB)"  
  
oPDF.end_page          ' finish page  
oPDF.close             ' close PDF document  
  
set oPDF = Nothing  
End Sub
```

Error handling in Visual Basic. A Visual Basic program can detect when an error happens, and react upon the error. Catching Exceptions in Visual Basic is achieved with an *On Error GoTo* clause:

```
Sub main()  
    Dim oPDF As PDFlib_com.PDF  
    On Error GoTo ErrExit
```

```
...some PDFlib instructions...
```

```
End
```

```
ErrExit:
```

```
MsgBox Hex(Err.Number) & ": " & Err.Description
```

```
End Sub
```

Note You can disable error handling in VBScript with the undocumented `On Error GoTo 0` statement (i.e., using zero as address for the `GoTo` statement).

Unicode support in Visual Basic. Visual Basic supports Unicode internally. (VB's program editor, however, doesn't seem to be fully Unicode-aware). Unicode strings can be constructed from numerical values using the `ChrW` function:

```
oPDF.set_info "nativeunicode", "true"
```

```
Unicodetext = ChrW(&H39B) & ChrW(&H39F) & ChrW(&H393) & ChrW(&H39F) & ChrW(&H3A3)
```

2.2.8 Using PDFlib with Windows Script Host

Windows Script Host¹ supports JavaScript and VBScript, the details of which have already been discussed in Section 2.2.6, »Using PDFlib with Active Server Pages«. For this reason, we will only present the hello example for VBScript here.

The »Hello world« example for Windows Script Host (WSH) with VBScript.

```
' hello.vbs
,
' PDFlib client: hello example for ActiveX with Windows Script Host and VBS
' Requires the PDFlib ActiveX component
,
Option Explicit
On Error Resume Next

Dim font
Dim oPDF

Set oPDF = CreateObject("PDFlib_com.PDF")

' Open new PDF file
if (oPDF.open_file("hello_ax_vbs.pdf") = -1) then
    WScript.Echo "Couldn't open PDF file!"
    WScript.Quit(1)
end if

oPDF.set_info "Creator", "hello.asp"
oPDF.set_info "Author", "Thomas Merz"
oPDF.set_info "Title", "Hello, world (Active X/VBS)!"

' start a new page
oPDF.begin_page 595, 842

font = oPDF.findfont("Helvetica-Bold", "winansi", 0)
```

¹ WSH is available in command-line (`cscript.exe`) and windowing flavors (`wscript.exe`). WSH is included in Microsoft Internet Explorer 5, Windows 98, and Windows 2000. For more information see <http://msdn.microsoft.com/scripting>.

```
oPDF.setFont font, 24
```

```
oPDF.set_text_pos 50, 700  
oPDF.show "Hello, world!"  
oPDF.continue_text "(says ActiveX/VBS)"  
oPDF.end_page  
oPDF.close  
set oPDF = Nothing
```

2.2.9 Using PDFlib with ColdFusion

Special considerations for ColdFusion¹. Allaire ColdFusion for Windows (but not ColdFusion Express) supports COM objects such as PDFlib via its *CFOBJECT* tag. After installing (and thereby registering) the PDFlib ActiveX edition no further steps are required in order to use PDFlib from ColdFusion. Since we couldn't find any way to generate PDF in memory and stream it to the client with ColdFusion the example below generates a temporary PDF file instead. According to the ColdFusion documentation functions without any parameters (e.g., *save*) should be called with an empty pair of parentheses.

Note PDFlib currently only supports the Windows version of ColdFusion.

The »Hello world« example for ColdFusion.

```
<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.0 Transitional//EN">  
<html>  
<head>  
  <title>PDFlib ColdFusion sample</title>  
</head>  
  
<body>  
  
<CFOBJECT type="COM" name="oPDF" class="PDFlib_com.PDF" action="CREATE">  
  
<CFSET tmp_file=GetTempFile("C:\\tmp\\", "hello_cfm.pdf")>  
<CFSET res=oPDF.open_file(tmp_file)>  
<CFIF res is -1>  
  Couldn't open PDF file!  
<CFABORT/>  
</CFIF>  
  
<CFSET oPDF.set_info("Creator", "hello.cfm")>  
<CFSET oPDF.set_info("Author", "Thomas Merz")>  
<CFSET oPDF.set_info("Title", "Hello, world (ActiveX/ColdFusion)")>  
  
<!-- start a new page -->  
<CFSET oPDF.begin_page(595, 842)>  
<CFSET font = oPDF.findfont("Helvetica-Bold", "winansi", 0)>  
  
<CFSET oPDF.setFont(font, 24)>  
  
<CFSET oPDF.set_text_pos(50, 700)>  
<CFSET oPDF.show("Hello, world!")>  
<CFSET oPDF.continue_text("(says ActiveX/ColdFusion)")>  
  
<CFSET oPDF.end_page()>
```

1. See <http://www.allaire.com/products/coldfusion>

```

<CFSET oPDF.close()>

<CFCONTENT TYPE="application/pdf" FILE="#tmp_file#" DELETEFILE="yes">

</body>
</html>

```

Error handling in ColdFusion. ColdFusion supports structured exception handling which can be used for dealing with PDFlib's COM exceptions. The error message supplied by PDFlib can be accessed via ColdFusion's *CFCATCH.detail* variable. In order to catch PDFlib exceptions from your ColdFusion code use the following basic structure:

```

<CFTRY>
...some PDFlib instructions...
<CFCATCH>
    <CFOUTPUT>
    PDFlib exception caught: #CFCATCH.detail#
    <CFABORT/>
    </CFOUTPUT>
</CFCATCH>
</CFTRY>

```

Unicode support in ColdFusion. Since we couldn't find any Unicode support in ColdFusion, PDFlib's automatic Unicode conversion features are not available in ColdFusion. ColdFusion developers must manually construct their Unicode strings according to Section 3.3.8, »Unicode Support«.

2.2.10 Using PDFlib with Borland Delphi

Special considerations for Borland Delphi¹. In order to use PDFlib with Delphi 5 you must reference the PDFlib type library in your project as follows: In *Project, Import Type Library...*, choose PDFlib from the list (or use *Add...* to select the *pdflib_com.dll* on disk if PDFlib is missing from the list), pick an arbitrary entry in *Palette page*, and *Install...* This must only be done once for each Delphi installation, not once for each project.

The »Hello world« example for Borland Delphi.

```

unit hello;

interface

uses Windows, Classes, Forms, Dialogs, ExtCtrls, Controls, StdCtrls,
    PDFlib_com_TLB, OleServer;

type
    TForm1 = class(TForm)
        Panel1: TPanel;
        InsertBtn: TButton;
        pdf: TPDF;
        procedure InsertBtnClick(Sender: TObject);
    end;

var
    Form1: TForm1;

```

1. See <http://www.borland.com/delphi>

implementation

```
{$R *.DFM}

procedure TForm1.InsertBtnClick(Sender: TObject);
var
  f: Integer;
begin
  if pdf.open_file('hello_delphi.pdf') = -1 then begin
    ShowMessage('Couldn't open hello_delphi.pdf!');
    Exit;
  end;

  pdf.set_info('Creator', 'hello.pas');
  pdf.set_info('Author', 'Rainer Schaaf');
  pdf.set_info('Title', 'Hello World (Delphi)');

  pdf.begin_page(595,842);
  f := pdf.findfont('Helvetica-Bold', 'host', 0);

  pdf.setfont(f, 18.0);
  pdf.set_text_pos(50, 700);
  pdf.show('Hello World');
  pdf.continue_text('(says Delphi)');
  pdf.end_page;
  pdf.close;
end;

end.
```

Error handling in Borland Delphi. Delphi supports structured exception handling which can be used for dealing with exceptions thrown from OLE objects such as PDFlib. In order to catch PDFlib exceptions from your Delphi code use the following basic structure:

```
uses SysUtils;

...

try
  ...some PDFlib instructions...
except
  On E: Exception do begin
    ShowMessage(E.Message);
  end;
end;
```

Unicode support in Borland Delphi. Delphi supports Unicode natively with its *WideString* data type. If you want to leverage PDFlib's Unicode support from Delphi you must use your strings of type *WideString*:

```
unicodetext: WideString;
...
pdf.set_parameter('nativeunicode', 'true');
unicodetext := #$039B;
```

2.3 C Binding

(This section is not included in this edition of the PDFlib manual.)

2.4 C++ Binding

(This section is not included in this edition of the PDFlib manual.)

2.5 Java Binding

(This section is not included in this edition of the PDFlib manual.)

2.6 Perl Binding

(This section is not included in this edition of the PDFlib manual.)

2.7 PHP Binding

(This section is not included in this edition of the PDFlib manual.)

2.8 Python Binding

(This section is not included in this edition of the PDFlib manual.)

2.9 RPG Binding

(This section is not included in this edition of the PDFlib manual.)

2.10 Tcl Binding

(This section is not included in this edition of the PDFlib manual.)

3 PDFlib Programming Concepts

3.1 General Programming Issues

3.1.1 The PDFlib Demo Stamp and Serial Numbers

All binary PDFlib and PDI versions 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 will display a `www.pdflib.com` demo stamp (the »nagger«) cross all generated pages. Companies which are seriously interested in PDFlib licensing and wish to get rid of the nagger during the evaluation phase or for prototype demos can submit their company and project details to `sales@pdflib.com`, and request a temporary serial string.

Once you purchased a PDFlib or PDI serial string you must apply it in order to get rid of the demo stamp. This can be achieved by supplying the serial at runtime:

```
oPDF.set_parameter("serial", "...your serial string...")
```

The serial string must be set only once, immediately after instantiating the PDFlib object (i.e., after `PDF_new()` or equivalent call). PDFlib and PDI serial strings are platform-specific, and must be purchased for a particular platform.

Users of the ActiveX edition can supply the serial string when they install PDFlib/PDI using the supplied installer. In this case the above function call is not required (see also Section 2.2.2, »Installing the PDFlib ActiveX Edition«).

Note that PDFlib and PDI are different products, and require different serial strings although they are delivered in a single package. PDI serials will also be valid for PDFlib, but not the other way round. Also, PDFlib and PDI serial strings are platform-dependent, and can only be used on the platform for which they have been purchased.

3.1.2 PDFlib Program Structure

PDFlib applications must obey certain structural rules which are very easy to understand. Writing applications according to these restrictions is straightforward. For example, you don't have to think about opening a page first before closing it. Since the PDFlib API is very closely modelled after the document/page paradigm, generating documents the »natural« way usually leads to well-formed PDFlib client programs.

PDFlib enforces correct ordering of function calls with a strict scoping system (see Section 4.1, »Data Types, Naming Conventions, and Scope«). The function descriptions document the allowed scope for a particular functions. Calling a function from a different scope will immediately trigger a PDFlib exception. PDFlib will also throw an exception if bad parameters are supplied by a library client.

3.1.3 Generating PDF Documents directly 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 one big chunk at the end (after `close()`). 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.

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

```
oPDF.open_file ("")
...create document...
oPDF.close

' Fetch the buffer with the PDF and write to a file
' This is rather pointless in VB but useful in ASP
Open "file.pdf" For Binary Access Write As #1
Put #1, , oPDF.get_buffer
Close #1
Set oPDF = Nothing
```

Note Fetching PDF data from a buffer requires binary access, and may not be usable from all environments due to restrictions of the respective development environment.

3.1.4 Error 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, then, is to use conventional error reporting mechanisms (read: special function return codes) for function calls which may go wrong often times, 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
- ▶ Using a font for which metrics information cannot be found
- ▶ Trying to open a corrupt image file
- ▶ Trying to import an encrypted PDF file

PDFlib signals such errors by returning a special value (usually `-1`) as documented in the API reference. 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 functions (e.g., trying to draw a circle with a negative radius)

If the library detects such an exceptional situation, a COM exception is thrown instead of passing special return values to the caller.

PDFlib exceptions fall into one of several categories as shown in Table 2.3. The error handler will receive the type of PDFlib error along with a descriptive message.

Non-fatal error messages (warnings) generally indicate some problem in your PDFlib code which you should investigate more closely. However, processing may continue in case of non-fatal errors. For this reason, you can suppress warnings using the following function call:

```
oPDF.set_parameter "warning", "false"
```

The suggested strategy is to enable warnings during the development cycle (and closely examine possible warnings), and disable warnings in a production system.

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 in PDF lingo) has the origin in the lower left corner of the page, and uses the DTP point as unit:

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

The first coordinate increases to the right, the second coordinate increases upward. 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 *rotate()*, *scale()*, *translate()*, and *skew()*. If the user space 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.

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 (take care to use suitable material since cheap overhead transparencies do not withstand heat, and may ruin your laser printer!) may provide a useful tool for preparing PDFlib development.

Don't be misled by PDF printouts which seem to experience wrong page dimensions. These may be wrong because of some common reasons:

- ▶ The *Fit to Page* (or *Shrink oversized pages to paper size*) option has been checked in Acrobat's print dialog, resulting in scaled print output.
- ▶ Non-PostScript printer drivers are not always able to retain the exact size of page objects.

Note Hypertext functions, such as those for creating text annotations, links, and file annotations are not affected by user space transformations, and always use the default coordinate system instead.

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:

```
oPDF.[scale] 28.3465, 28.3465
```

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

Rotating objects. It is important to understand that object 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 future objects.

The following example generates some horizontal text, and rotates the coordinate system in order to show vertical text. The save/restore nesting makes it easy to continue with vertical text in the original coordinate system after the vertical text is done:

```
oPDF.set_text_pos 50, 600
oPDF.show "This is horizontal text"
textx = oPDF.get_value("textx", 0)           ' determine text position
texty = oPDF.get_value("texty", 0)           ' determine text position

oPDF.save
    oPDF.translate textx, texty               ' move origin to end of text
    oPDF.rotate 90                           ' rotate coordinates
    oPDF.set_text_pos 18, 0                  ' provide for distance from horizontal text
    oPDF.show "vertical text"
oPDF.restore

oPDF.continue_text "horizontal text continues"
```

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 additional calls are required in order to avoid text being displayed in a mirrored sense. In order to set up a coordinate system with the origin in the top left corner of the page and the y -coordinate pointing downwards while maintaining the usual text direction (text stands upright on the page) use the following code sequence:

```
oPDF.begin_page Width, Height                ' set up the page dimensions
oPDF.translate 0, Height                     ' move the coordinate origin
oPDF.[Scale] 1, -1                          ' reflect at the horiz. axis

font = oPDF.findfont("Helvetica-Bold", "host", 0) ' sample text
oPDF.setfont font, -18                      ' make the text point upwards
oPDF.set_value "horizscaling", -100         ' compensate for the mirroring

oPDF.set_text_pos 50, 100                   ' now use top-down coordinates
oPDF.show "Hello world!"
```

In order to format text into a text box with the upper right corner at (x, y) , width w , and height h use the following idiom (this is required because the function adds h to the starting y position):

```
c = show_boxed(text, x, y-h, w, h, "justify", "")
```

Similarly, the following idiom can be used in order to correctly place images when using top-down coordinates:

```
' Place the image in the lower left corner of the page
oPDF.save
    oPDF.translate 0, Height ' temporarily translate origin to lower left corner
    oPDF.scale 1, -1
    oPDF.place_image lImage, 0, 0, 1
oPDF.restore
```

3.2.2 Page and Coordinate Limits

Page sizes. Although PDF and PDFlib don't impose any restrictions on the usable page size, Acrobat implementations suffer from architectural limits regarding the page size. Note that other PDF interpreters may well be able to deal with larger or smaller document formats. If run in Acrobat 3 compatibility mode PDFlib will throw a *RuntimeError* exception if the Acrobat 3 limits are exceeded; if run in Acrobat 4 (the default) or 5 compatibility mode and the Acrobat 4 limits are exceeded, PDFlib will only issue a non-fatal warning message. Common standard page size dimensions can be found in Table 3.1.

Table 3.1. Minimum and maximum page size of several PDF consumers

PDF viewer	minimum page size	maximum page size
Acrobat 3	1" = 72 pt = 2.54 cm	45" = 3240 pt = 114.3 cm
Acrobat 4 and 5	1/24" = 3 pt = 0.106 cm	200" = 14400 pt = 508 cm

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 the box entries of Acrobat 4/PDF 1.3. 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.
- ▶ **CropBox:** the region to which the page contents are to be clipped;
- ▶ **TrimBox:** the intended dimensions of the finished page after trimming;
- ▶ **ArtBox:** extent of the page's meaningful content;
- ▶ **BleedBox:** the region to which the page contents are to be clipped when output in a production environment.

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.

Number of pages in a document. There is no intrinsic limit in PDFlib regarding the number of generated pages in a document. While previous versions of PDFlib generated output which resulted in bad Acrobat performance when navigating large files, PDFlib 4 introduces an improvement in the generated PDF structures which significantly accelerates document navigation in Acrobat even for documents with hundreds of thousands of pages.

Output accuracy and coordinate range. PDFlib's numerical output accuracy has been carefully chosen to match the requirements of PDF and the supported environments, while at the same timing minimizing output file size. As detailed in Table 3.2 PDFlib's accuracy depends on the absolute value of coordinates. While most developers may

safely ignore this issue, demanding applications should take care in their scaling operations in order to not exceed PDF's built-in coordinate limits.

Table 3.2. Output accuracy and coordinate range

<i>absolute value</i>	<i>output</i>
$0 \dots 0.000015$	0
$0.000015 \dots 32767.999999$	rounded to four decimal digits, configurable for up to six digits
$32768 \dots 2^{31} - 1$	rounded to next integer
$\geq 2^{31}$	an exception of type <i>ValueError</i> will be raised

3.2.3 Paths and Color

Graphics paths. A path is a shape made of an arbitrary number of straight lines, rectangles, or curves. A path may consist of several disconnected sections, called subpaths. There are several operations which can be applied to a path (see Section 4.4.4, »Path Painting and Clipping«):

- ▶ Stroking draws a line along the path, using client-supplied parameters for drawing.
- ▶ Filling paints the entire region enclosed by the path, using client-supplied parameters for filling.
- ▶ Clipping reduces the imageable area for subsequent drawing operations by replacing the current clipping area (which is the page size by default) with the intersection of the current clipping area and 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 required.

It is an error to construct a path without applying one of the above operations on it. PDFlib's scoping system ensures that clients obey to this restriction. These rules may easily 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:

```
oPDF.moveto 100, 100
oPDF.lineto 200, 100
oPDF.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.

Color. PDFlib clients may specify the colors used for filling and stroking the interior of paths and text characters. Colors may be specified in one of several color spaces:

- ▶ gray values between 0=black and 1=white;
- ▶ RGB triples, i.e., three values between 0 and 1 specifying the percentage of red, green, and blue; (0, 0, 0)=black, (1, 1, 1)=white;
- ▶ four CMYK values between 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 this is different from the RGB specification.
- ▶ spot color: an arbitrarily named color with an alternate representation in one of the other color spaces above; this is generally used for preparing documents which are intended to be printed on an offset printing machine with one or more custom col-

ors. The tint value (percentage) ranges from 0=no color to 1=maximum intensity of the spot color.

- ▶ pattern: tiling with an object composed of arbitrary text, vector, or image graphics (patterns are not supported in Acrobat 3 compatibility mode since they don't show up on screen with Acrobat 3).

The default value for stroke and fill color is black, i.e. (0, 0, 0) in the RGB color space.

3.2.4 Templates

Templates in PDF. 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 acting on a regular page). After the template is finished it can be used much 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 suggest themselves 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. Other typical examples for template usage include crop and registration marks or custom Asian glyphs.

Note PDF templates are an efficient means for saving space in a PDF file. However, this advantage is usually not retained when printing template-based PDF files to a PostScript printer. Depending on the number of templates used, you should be prepared for print jobs which are significantly larger than the corresponding PDF files.

Using templates with PDFlib. Templates can only be *defined* outside of a page description, and can be *used* within a page description. However, templates may also contain other templates. Obviously, using a template within its own definition is not possible. Referring to an already defined template on a page is achieved with the *place_image()* function just like images are placed on the page (see Section 3.4.2, »Code Fragments for Common Image Tasks«). The general template idiom in PDFlib looks as follows:

```
' define the template
lTemplate = oPDF.begin_template(template_width, template_height);
...place marks on the template using text, vector, and image functions...
oPDF.end_template
...
oPDF.begin_page page_width, page_height
' use the template
oPDF.place_image lTemplate, 0.0, 0.0, 1.0
...more page marking operations...
oPDF.end_page
...
oPDF.close_image lTemplate
```

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

- ▶ The functions in Section 4.6, »Image Functions«, except *place_image()* and *close_image()*. This is not a big restriction since images can be opened outside of a template definition, and freely be used within a template (but not opened).
- ▶ The functions in Section 4.8, »Hypertext Functions«. Hypertext elements must always be defined on the page where they should appear in the document, and cannot be generated as part of a template.

Note You can apply all image manipulation algorithms in Section 3.4.2, »Code Fragments for Common Image Tasks« to templates, too. Simply substitute the template width for *get_value("imagewidth", image)*, similar for template and image height.

Templates and the graphics state. When a template is placed on a page, it will inherit all graphics state parameters from the page unless these are explicitly set within the template definition. For example, if a page description sets the current color to red and places a template which doesn't explicitly set the current color, text and vector elements on the template will be drawn in red color, too. This behavior can be used to change the color of templates, but may also be undesirable for some applications. You have the following choices in this situation:

- ▶ Templates which will always be drawn in the same color should specify all required graphics state parameters within the template definition.
- ▶ Templates which will be recolored (more precisely: which shall inherit graphics state parameters from the surrounding page) should not set any graphics state parameter. Using such context-dependent templates requires appropriate setup on the surrounding page.
- ▶ If the behavior of a template is not known (especially when the template consists of an imported PDF page of unknown origin) all affected graphics state parameters should be reset to their defaults. Since it is usually not known which parameters are used within an imported PDF page, using *initgraphics()* is the easiest way to assure correct behavior. Note that this function also resets the current transformation matrix.

In a similar fashion, some property of the graphics state may be modified after placing a template or imported PDI page. It is safer to explicitly set the color etc. after placing a template or PDI page, or to bracket the template or page with *save()/restore()*.

Template support in third-party software. Templates (form XObjects) are an integral part of the PDF specification, and can be perfectly viewed and printed with Acrobat. However, since this type of PDF construct is rarely generated by Acrobat Distiller, not all PDF consumers are prepared to deal with it. For example, not even Acrobat 4's touch-up tool can be used for manipulating templates (this has been fixed in Acrobat 5). Similarly, the PitStop 4.5 PDF editor can only move templates, but cannot access individual elements within a template. On the other hand, Adobe Illustrator 9 fully supports templates.

3.3 Text Handling

3.3.1 The PDF Core Fonts

PDF viewers support a core set of 14 fonts which need not be embedded in any PDF file. Even when a font isn't embedded in the PDF file, PDF and therefore PDFlib need to know

about the width of individual characters. For this reason, metrics information for the core fonts is already built into the PDFlib binary. However, the builtin metrics information is only available for the native host encoding (see below). Using another encoding than the host encoding requires metrics information files. Metrics files for the PDF core fonts are included in the PDFlib distribution in order to make it possible to use encodings other than the host encoding. The core fonts are the following:

*Courier, Courier-Bold, Courier-Oblique, Courier-BoldOblique,
Helvetica, Helvetica-Bold, Helvetica-Oblique, Helvetica-BoldOblique,
Times-Roman, Times-Bold, Times-Italic, Times-BoldItalic,
Symbol, ZapfDingbats*

3.3.2 8-Bit Encodings built into PDFlib

PDF supports flexible text encodings (the mapping of numerical code values to character glyphs) for 8-bit text fonts. PDFlib includes provisions for supporting diverse encoding vectors for dealing with text. The builtin encoding vectors are referred to via symbolic names. Table 3.3 lists the symbolic encoding names supported internally by PDFlib. Additional encodings are available in external encoding files distributed with PDFlib (see below), or can be defined by the user (see Section 3.3.3, »Custom Encoding and Code Page Files for 8-Bit Encodings«). All supported encodings can be arbitrarily mixed in one document. You may even use different encodings for a single font, although the need to do so will only rarely arise.

Note Not all encodings can be used with a given font. The user is responsible for making sure that the font contains all characters required by a particular encoding. This can even be problematic with Acrobat's core fonts.

Table 3.3. Builtin character encodings supported by PDFlib

encoding	description
<i>winansi</i>	<i>Windows code page 1252, a superset of ISO 8859-1</i>
<i>macroman</i>	<i>Mac Roman encoding, i.e., the default Macintosh character set</i>
<i>ebcdic</i>	<i>EBCDIC code page 1047 as used on IBM AS/400 and S/390 systems</i>
<i>builtin</i>	<i>Original encoding used by non-text (symbol) or non-Latin text fonts</i>
<i>host</i>	<i>macroman on the Mac, ebcdic on EBCDIC-based systems, and winansi on all others</i>

The winansi encoding. This encoding reflects the Windows ANSI character set, more specifically code page 1252 including the three characters which Microsoft added for Windows 98 and Windows 2000 (*Euro*, *Zcaron*, and *zcaron*). The *winansi* encoding is a superset of ISO 8859-1 (Latin-1) and can therefore also be used on Unix systems.

Note Most PostScript fonts do not yet contain the three additional Windows characters. They are supported by the core fonts in Acrobat 4, however.

The macroman encoding. This encoding reflects the MacOS character set, albeit with the old currency symbol at position 219, and not the Euro character as redefined by Apple (this incompatibility is dictated by the PDF specification). Also, this encoding does not include the Apple glyph and the mathematical symbols as defined in the MacOS character set.

The ebcdic encoding. This encoding relates to the EBCDIC (*Extended Binary Coded Decimal Interchange Code*) defined by IBM and used on the IBM AS/400, S/390, and other midrange and mainframe systems. More specifically, PDFlib's *ebcdic* encoding uses the EBCDIC code page 1047. As with all other PDFlib encodings, *ebcdic* encoding is always available for generating PDF output, and not only on native EBCDIC machines. The difference, however, is that on those machines the built-in metrics for the core fonts are sorted according to *ebcdic* encoding, and that *host* encoding (see below) also relates to *ebcdic* encoding.

The builtin encoding. The encoding name *builtin* doesn't describe a particular character ordering but rather means »take this font as it is, and don't mess around with the character set«. This concept is sometimes called a »font specific« encoding and is very important when it comes to non-text fonts (such as logo and symbol fonts), or non-Latin text fonts (such as Greek and Cyrillic). Such fonts cannot be reencoded using one of the supported encodings since their character names don't match those in these encodings. Therefore, *builtin* must be used for all symbolic or non-text fonts, such as Symbol and ZapfDingbats. Non-text fonts can be recognized by the following entry in their AFM file:

```
EncodingScheme FontSpecific
```

Text fonts can be reencoded (adjusted to a certain code page or character set), while symbolic fonts can't, and must use *builtin* encoding instead.

Note Unfortunately, many typographers and font vendors didn't fully grasp the concept of font specific encodings (this may be due to less-than-perfect production tools). For this reason, there are many Latin text fonts labeled as FontSpecific encoding, and many symbol fonts incorrectly labeled as text fonts.

The host encoding. Like *builtin*, the *host* encoding plays a special role since it doesn't refer to some fixed character set. Instead, *host* encoding will be mapped to *macroman* on the Mac, *ebcdic* on EBCDIC-based systems, and *winansi* on all others. The *host* encoding is primarily useful as a vehicle for writing platform-independent test programs (like those contained in the PDFlib distribution) or other encoding-wise simple applications. Assuming that PDFlib client programs are always encoded in the host's native encoding, such programs will always generate PDF text output with the »correct« encoding. Contrary to all other aspects of PDFlib, the concept of a *host* encoding is inherently non-portable. For this reason *host* encoding is not recommended for production use.

3.3.3 Custom Encoding and Code Page Files for 8-Bit Encodings

In addition to a number of predefined encodings (see Section 3.3.2, »8-Bit Encodings built into PDFlib«) PDFlib supports user-defined 8-bit encodings in order to make PDFlib's font handling even more flexible. User-defined encodings are the way to go if you want to deal with some character set which is not internally available in PDFlib, such as EBCDIC character sets different from the one supported internally in PDFlib. In addition to encoding tables defined by PostScript glyph names PDFlib also accepts code page tables which describe a mapping from Unicode to a set of up to 256 characters. These characters can be accessed with 8-bit character codes.

The following tasks must be done before a user-defined encoding can be leveraged in a PDFlib program:

- ▶ Generate a description of the encoding in a simple text format.
- ▶ Configure the encoding in the PDFlib resource file (see Section 3.3.6, »Resource Configuration and the UPR Resource File«) or via `set_parameter()`.
- ▶ Provide a font (metrics and possibly outline file) that supports all characters used in the encoding. Of course, the characters in the font must use the correct PostScript glyph names as defined in the encoding table.

The encoding file simply lists glyph names and numbers line by line. As an example, the following excerpt shows the encoding definition for the ISO 8859-2 (Latin 2) character set:

```
% Encoding definition for PDFlib
% ISO 8859-2 (Latin-2)
space          32      % 0x20
exclam        33      % 0x21
quotedbl      34      % 0x22
...more glyph assignments...
yacute        253     % 0xFD
tcommaaccent  254     % 0xFE
dotaccent     255     % 0xFF
```

The next example shows a snippet from a Unicode code page for the same ISO 8859-2 character set:

```
% Code page definition for PDFlib
% ISO 8859-2 (Latin-2)
0x0020        32      % 0x20
0x0021        33      % 0x21
0x0022        34      % 0x22
...more glyph assignments...
0x00FD        253     % 0xFD
0x0163        254     % 0xFE
0x02D9        255     % 0xFF
```

More formally, the contents of an encoding or code page file are governed by the following rules:

- ▶ Comments are introduced by a percent '%' character, and terminated by the end of the line.
- ▶ The first entry in each line is either a PostScript character name or a hexadecimal Unicode value composed of a `0x` prefix and four hex digits (upper or lower case). This is followed by whitespace and a hexadecimal or decimal character code in the decimal range 0–255. Only Unicode values in the Adobe Glyph List (AGL) are allowed (see below).
- ▶ Character codes which are not mentioned in the encoding file are assumed to be undefined. Alternatively, a Unicode value of `0x0000` or the character name `.notdef` can be used for unencoded characters.

As a naming convention we refer to name-based tables as encoding files (**.enc*), and Unicode-based tables as code page files (**.cpg*), although actually PDFlib treats both kinds in the same way (and doesn't care about file names). In fact, PDFlib will automatically convert between name-based encoding files and Unicode-based code page files whenever it is necessary. This conversion is based on Adobe's standard list of PostScript glyph

names (the Adobe Glyph List, or AGL¹), which is built into PDFlib. Encoding files are required for PostScript fonts with non-standard glyph names, while code pages are more convenient when dealing with Unicode-based TrueType fonts.

The relationship between the name of the encoding file and the name of the actual encoding (to be used with `findfont()`) is specified in PDFlib's resource file or via `set_parameter()` (see Section 3.3.6, »Resource Configuration and the UPR Resource File«).

Distributed encoding files. The PDFlib distribution contains several encoding and code page files (see Table 3.4) which may be useful if you need to use one of the supplied encodings directly, or want to use it as a starting point for writing your own encoding files. In order to use these, the PDFlib resource configuration file and font metrics files must be accessible (see Section 3.3.6, »Resource Configuration and the UPR Resource File«).

Table 3.4. Additional external character encodings distributed with PDFlib

encoding	description
iso8859-2	Latin-2 supports the Slavic languages of Central Europe which use the Latin alphabet. Acrobat 4's core fonts do not contain all characters for ISO 8859-2.
iso8859-3	Latin-3 covers Esperanto and Maltese.
iso8859-4	Latin-4 covers Estonian, the Baltic languages, Greenlandic, and Lappish.
iso8859-5	ISO 8859-5 (Cyrillic) covers Bulgarian, Russian, and Serbian. Acrobat 4 does not correctly implement TrueType handling for Cyrillic characters.
iso8859-6	ISO 8859-6 covers Arabic, but not Persian or Pakistani/Urdu.
iso8859-7	ISO 8859-7 covers modern Greek
iso8859-8	ISO 8859-8 covers Hebrew and Yiddish.
iso8859-9	Latin-5 (yes, that's 5, not 9!) supports Danish, Dutch, English, Finnish, French, German, Irish, Italian, Norwegian, Portuguese, Spanish, Swedish, and Turkish. In addition, this PDFlib encoding contains the characters 130-159 (&H82-&H9F) as defined in the Windows code page 1254 (Turkish). Acrobat 4's core fonts do not contain the following characters for ISO 8859-9: Gbreve, gbreve, Idotaccent, Scommaaccent, scommaaccent.
iso8859-10	Latin-6 is a variation of Latin-4 and covers the Nordic area.
iso8859-15	Latin-9: this character set is a variation of Latin-1 which adds the Euro character as well as some missing French and Finnish characters. Latin-9 is sometimes also dubbed Latin-o, although this is not the official name.
cp1250	Windows code page 1250 (Central European). Acrobat 4's core fonts do not contain all characters for code page 1250.
cp1251	Windows code page 1251 (Cyrillic). Acrobat 4 does not correctly implement TrueType handling for Cyrillic characters.
cp1253	Windows code page 1253 (Greek)
cp1254	Windows code page 1254 (Turkish)
cp1255	Windows code page 1255 (Hebrew)
cp1256	Windows code page 1256 (Arabic)
cp1257	Windows code page 1257 (Baltic)
cp1258	Windows code page 1258 (Viet Nam)

1. The AGL can be found at <http://partners.adobe.com/asn/developer/type/glyphlist.txt>

Finding PostScript character names. In order to write a custom encoding file or find fonts which can be used with one of the supplied encodings you will have to find information about the exact definition of the character set to be defined by the encoding, as well as the exact glyph names used in the font files. You must also ensure that a chosen font provides all necessary characters for the encoding. For example, the core fonts supplied with Acrobat 4 do not support ISO 8859-2 (Latin 2) nor Windows code page 1250. If you happen to have the FontLab¹ font editor (by the way, a great tool for dealing with all kinds of font and encoding issues), you may use it to find out about the encodings supported by a given font (look for »code pages« in the FontLab documentation).²

For the convenience of PDFlib users, the PostScript program *print_glyphs.ps* in the distribution fileset can be used to find the names of all characters contained in a PostScript font. In order to use it, enter the name of the font at the end of the PostScript file and send it (along with the font) to a PostScript Level 2 or 3 printer, or view it with a Level-2-compatible PostScript viewer such as Ghostscript³. The program will print all characters in the font, sorted alphabetically by glyph name.

If a font does not contain a character required for a custom encoding, it will be missing in the PDF document.

3.3.4 Hypertext Encoding

PDF supports two methods for encoding hypertext elements such as bookmarks, annotations, and document information fields. Up to Acrobat 3, all hypertext strings had to be encoded with a special 8-bit encoding called PDFDocEncoding (PDFDocEncoding can not be used for text used on page descriptions). Starting with Acrobat 4, Unicode strings can be used for all hypertext elements. For more information on Unicode see Section 3.3.8, »Unicode Support«.

PDFDocEncoding (see Figure 3.1) is a superset of ISO 8859-1 (Latin 1) and therefore contains all ASCII characters in the lower part. Although PDFDocEncoding and the Windows code page 1252 are quite similar, they differ substantially in the character range 128-160 (&H80-&HA0).

Many clients will be able to directly use PDFDocEncoding. However, since the Mac encoding substantially differs from PDFDocEncoding, it is necessary to convert Mac strings to PDFDocEncoding when it comes to hypertext elements, and non-ASCII special characters are to be used. Mac special characters must be converted to Unicode before they can be used in hypertext elements. This conversion must be performed by the client.

3.3.5 PostScript and TrueType Fonts

Font embedding in PDF. PDF supports fonts outside the set of 14 core fonts in several ways. PDFlib is capable of embedding font descriptions into the generated PDF output. Alternatively, a font descriptor consisting of the character metrics and some general information about the font (without the actual character outline data) can be embedded.

1. See <http://www.fontlab.com>

2. Useful raw material for writing encoding tables for a variety of standards and vendor-specific character sets can be found at <ftp://ftp.unicode.org/Public/MAPPINGS>; Information about the glyph names used in PostScript fonts can be found at <http://partners.adobe.com/asn/developer/typeforum/unicodegn.html> (although font vendors are not required to follow these glyph naming recommendations).

3. See <http://www.cs.wisc.edu/~ghost>

	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
0																
1	000	001	002	003	004	005	006	007	010	011	012	013	014	015	016	017
2	020	021	022	023	024	025	026	027	030	031	032	033	034	035	036	037
3	040	041	042	043	044	045	046	047	050	051	052	053	054	055	056	057
4	060	061	062	063	064	065	066	067	070	071	072	073	074	075	076	077
5	100	101	102	103	104	105	106	107	110	111	112	113	114	115	116	117
6	120	121	122	123	124	125	126	127	130	131	132	133	134	135	136	137
7	140	141	142	143	144	145	146	147	150	151	152	153	154	155	156	157
8	160	161	162	163	164	165	166	167	170	171	172	173	174	175	176	177
9	200	201	202	203	204	205	206	207	210	211	212	213	214	215	216	217
A	220	221	222	223	224	225	226	227	230	231	232	233	234	235	236	237
B	240	241	242	243	244	245	246	247	250	251	252	253	254	255	256	257
C	260	261	262	263	264	265	266	267	270	271	272	273	274	275	276	277
D	300	301	302	303	304	305	306	307	310	311	312	313	314	315	316	317
E	320	321	322	323	324	325	326	327	330	331	332	333	334	335	336	337
F	340	341	342	343	344	345	346	347	350	351	352	353	354	355	356	357
	360	361	362	363	364	365	366	367	370	371	372	373	374	375	376	377

Fig. 3.1 The PDFDocEncoding character set as defined in PDF 1.3 with hex and octal codes. Note the Euro character at position hexadecimal Ao = octal 240.

If a font is not embedded in a PDF document, Acrobat will take it from the target system if available, or construct a substitute font according to the font descriptor in the PDF. Table 3.5 lists different situations with respect to font usage, each of which poses different requirements on the necessary font and metrics files.

When a font with font-specific encoding (a symbol font) is used, but not embedded in the PDF output, the resulting PDF will be unusable unless the font in question is already natively installed on the target system (since Acrobat can only simulate Latin text fonts). Such PDF files are inherently nonportable, although they may be of use in controlled environments, such as intra-corporate document exchange.

PostScript fonts. PDFlib supports the following formats for PostScript metrics and outline data on all platforms:

- ▶ The platform-independent AFM (Adobe Font Metrics) and the Windows-specific PFM (Printer Font Metrics) format for metrics information. Since PFM files do not describe the full character metrics but only the glyphs used in Windows (code page 1252), they can only be used for the *winansi* or *builtin* encodings, while AFM-based font metrics can be rearranged to any encoding supported by the font.

Table 3.5. Different font usage situations and required metrics and outline files

font usage	font metrics file required?	font outline file required?
One of the 14 core fonts with PDFlib's host encoding ^{1,2}	no	no
One of the 14 core fonts with an encoding other than PDFlib's host encoding ²	yes (AFM files supplied with the PDFlib distribution)	no
Non-core PostScript fonts without embedding	yes	no
Non-core PostScript fonts with embedding	yes	yes
Additional font/encoding combinations for which the metrics have been compiled into PDFlib (see below)	no	yes, if embedding is requested
TrueType fonts with or without embedding	no	yes
Standard CID fonts ³	no	no
Non-standard CID fonts	(not supported)	(not supported)

1. See Section 3.3.1, »The PDF Core Fonts« for a list of core fonts.

2. See Section 3.3.2, »8-Bit Encodings built into PDFlib« for the definition of PDFlib's host encoding.

3. See Section 3.3.7, »CID Font Support for Japanese, Chinese, and Korean Text« for more information on CID fonts.

- ▶ The platform-independent PFA (Printer Font ASCII) and the Windows-specific PFB (Printer Font Binary) format for font outline information in the PostScript Type 1 format, (sometimes also called »ATM fonts«). PostScript Type 3 fonts are not supported.

If you can get hold of a PostScript font file, but not the corresponding metrics file, you can try to generate the missing metrics using one of several freely available utilities. For example, the T1lib package¹ contains the *type1afm* utility for generating AFM metrics from PFA or PFB font files.

PostScript font names. It is important to use the exact (case-sensitive) PostScript font name whenever a font is referenced in PDFlib. There are several possibilities to find a PostScript font's exact name:

- ▶ Open the font outline file (*.pfa or *.pfb), and look for the string after the entry /FontName. Omit the leading / character from this entry, and use the remainder as the font name.
- ▶ If you have ATM (Adobe Type Manager) installed, you can double-click the font (*.pfb) or metrics (*.pfm) file, and will see a font sample along with the PostScript name of the font.
- ▶ Open the AFM metrics file and look for the string after the entry FontName.

Note The PostScript font name may differ substantially from the Windows font menu name, e.g. »AvantGarde-Demi« (PostScript name) vs. »AvantGarde, Bold« (Windows font menu name). Also, the font name as given in any Windows .inf file is not relevant for use with PDF.

Performance notes for PostScript fonts. It is important to be aware of the impact of font handling issues on PDFlib's performance. Generally, the font metrics (either in-core or on file) are accessed whenever a certain font/encoding combination is used for the first time. Subsequent requests for the same combination will be satisfied from PDFlib's internal font cache without any further performance penalty. Regarding font handling performance, the following observations may be useful:

1. See <http://www.neuroinformatik.ruhr-uni-bochum.de/ini/PEOPLE/rmz/t1lib/t1lib.html>

- ▶ Due to their small size and binary nature, PFM metrics files can be read much faster than the text-based AFM metrics files. However, they cannot be used for arbitrary encodings.
- ▶ AFM files contain much useful information about many aspects of font usage, and can be used for arbitrary encodings. However, although only the bare character metrics are required for PDFlib, the complete AFM file must be parsed in a time-consuming manner. For performance-critical applications it might be worthwhile to strip the unneeded data (e.g., the kerning information) from the AFM file.

TrueType fonts. PDFlib supports TrueType and OpenType fonts on all platforms. The TrueType font file must be supplied in Windows TTF format (Macintosh resource format is not supported). OpenType fonts are only supported if they contain TrueType outlines (as opposed to OpenType fonts which contain PostScript outlines, also known as CFF fonts). Contrary to PostScript fonts, TrueType fonts do not require any additional metrics file since the metrics information can be extracted from the font file itself. PDFlib currently supports the following flavors of TrueType fonts:

- ▶ Standard latin text fonts with the Windows character set (in TrueType lingo: *cmap* table with platform id 3, encoding id 1). These must be used with PDFlib encoding *winansi*. Note that most (but not all) TrueType fonts for Windows can also be used with PDFlib encoding *macroman* since they contain the necessary Mac information (*cmap* table with platform id 1, encoding id 0) in addition to the Windows information.
- ▶ Unicode-compatible TrueType fonts. These can be used with any PDFlib encoding as long as the font actually contains the characters required by that encoding. You can check which code pages are supported by a particular font with the »font properties extension« mentioned below.
- ▶ Symbol fonts with a custom character set (in TrueType lingo: *cmap* table with platform id 3, encoding id 0). These must be used with PDFlib encoding *builtin*.

The above distinction between text and symbol fonts may seem obvious, but in practise it may be hard to find the appropriate category for a given font. For various reasons, text fonts may be coded as TrueType *symbol* fonts, and vice versa. In case of an encoding mismatch PDFlib tries to help, and supplies encoding suggestions in the message which is part of a PDFlib exception.

TrueType host fonts. In addition to accessing font files which have been configured via PDFlib's resource and parameter machinery (see Section 3.3.6, »Resource Configuration and the UPR Resource File«) TrueType fonts can also be fetched directly from the operating system. We refer to such fonts as *host fonts*. Instead of fiddling with font and configuration files simply install the font in the operating system (read: drop it into the *fonts* directory), and PDFlib will happily use it.

Note Host font support is restricted to TrueType fonts. Host fonts will not be used for embedding one of the core fonts.

TrueType font names. It is important to specify the exact (case-sensitive) TrueType font name whenever a font is referenced in PDFlib. This must be the Windows name of the font as it is exposed at the user interface. You can easily find this name by double-clicking the TrueType font file in Windows, and taking note of the full font name which will be displayed in the first line of the resulting window (without the *TrueType* or *Open-*

Type term in parentheses, of course). Do not use the entry in the second line after the label *Typeface name!* Also, 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, and must also be used in translated form in PDFlib.

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 differ from its genuine TrueType name (e.g., *TimesNewRomanPSMT* vs. *Times New Roman*).

If you want to examine TrueType fonts in more detail take a look at Microsoft's free »font properties extension«¹ which will display many entries of the font's TrueType tables in human-readable form.

Note Contrary to PostScript fonts, TrueType font names may contain blank characters (and often do).

Legal aspects of font embedding. It's important to note that mere possession of a font file may not justify embedding the font in PDF, even for holders of a legal font license. Many font vendors restrict embedding of their fonts. Some type foundries completely forbid PDF font embedding, others offer special online or embedding licenses for their fonts, while still others allow font embedding provided the fonts are subsetted. Please check the legal implications of font embedding before attempting to embed fonts with PDFlib. PDFlib will honour embedding restrictions which may be specified in a TrueType font. If the embedding flag in a TrueType font is set to *no embedding*², PDFlib will honor the font vendor's request, and reject any attempt at embedding the font.

Note PDFlib currently doesn't implement font subsetting.

3.3.6 Resource Configuration and the UPR Resource File

In order to make PDFlib's font and encoding handling platform-independent and customizable, a configuration file can be supplied for describing the available fonts along with the names of their outline and metrics files, and the names of additional encoding files. In addition to the static configuration file, dynamic configuration can be accomplished at runtime by adding resources with *set_parameter()*. For the configuration file we dug out a simple text format called *Unix PostScript Resource* (UPR) which came to life in the era of Display PostScript and is still in use on several systems. However, we will take the liberty of extending the original UPR format for our purposes. The UPR file format as used by PDFlib will be described below.³ There is a utility called *makepsres* (often distributed as part of the X Window System) which can be used to automatically generate UPR files from PostScript font outline and metrics files.

Note As an alternative to configuring fonts in UPR files or via *set_parameter()* you can make use of PDFlib host font feature in certain situations (see Section 3.3.5, »PostScript and TrueType Fonts«)

1. See <http://www.microsoft.com/typography/property/property.htm>

2. More specifically: if the *fsType* flag in the *OS/2* table of the font has a value of 2.

3. The complete specification can be found in the book »Programming the Display PostScript System with X« (Appendix A), available at <http://partners.adobe.com/asn/developer/PDFS/TN/DPS.refmanuals.TK.pdf>

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 255 characters.
- ▶ A backslash '\ ' escapes any character, including newline characters. This may be used to extend lines. Windows directory names must be separated by double backslashes '\\ ' or a single forward slash '/ '.
- ▶ The period character '.' serves as a section terminator, and must therefore be escaped when used at the start of any other line.
- ▶ All entries are case-sensitive.
- ▶ Comment lines may be introduced with a percent '%' character, and terminated by the end of the line.
- ▶ Whitespace is ignored everywhere.

UPR files consist of the following components:

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

```
PS-Resources-1.0
```

- ▶ A section listing all types of 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. This section exists for compatibility only, and is ignored by PDFlib.

Note Some PDFlib editions support a directory prefix entry in the UPR file. Do not use the directory prefix for the PDFlib ActiveX component since it substitutes the UPR prefix mechanism with Windows registry entries, and expects all files in the PDFlib fonts directory by default. If you want to access resources in other directories use double equal signs as described below.

- ▶ 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), an equal sign, and the corresponding relative or absolute file name for the resource. Relative file names will be interpreted relative to PDFlib's *fonts* directory. Using a double equal sign forces the file name to be interpreted absolute, i.e., the prefix is not used.

Supported resource categories. The resource categories supported by PDFlib are listed in Table 3.6. Other resource categories may be present in the UPR file for compatibility with Display PostScript installations, but they will silently be ignored.

Table 3.6. Resource categories supported in PDFlib

resource category name	explanation
FontAFM	PostScript font metrics file in AFM format
FontPFM	PostScript font metrics file in PFM format
FontOutline	PostScript, TrueType or OpenType font outline file
Encoding	text file containing an 8-bit encoding or code page table

Redundant resource entries should be avoided. For example, do not include multiple entries for a certain font's metrics data. Also, the font name as configured in the UPR file should exactly match the actual font name in order to avoid confusion (although PDFlib does not enforce this restriction).

Sample UPR file. The following listing gives an example of a UPR configuration file as used by PDFlib. It describes the 14 PDF core fonts' metrics, plus metrics and outline files for some additional fonts, plus a custom encoding:

```
PS-Resources-1.0
FontAFM
FontPFM
FontOutline
Encoding
.
FontAFM
Code-128=Code_128.afm
Courier=Courier.afm
Courier-Bold=Courier-Bold.afm
Courier-BoldOblique=Courier-BoldOblique.afm
Courier-Oblique=Courier-Oblique.afm
Helvetica=Helvetica.afm
Helvetica-Bold=Helvetica-Bold.afm
Helvetica-BoldOblique=Helvetica-BoldOblique.afm
Helvetica-Oblique=Helvetica-Oblique.afm
Symbol=Symbol.afm
Times-Bold=Times-Bold.afm
Times-BoldItalic=Times-BoldItalic.afm
Times-Italic=Times-Italic.afm
Times-Roman=Times-Roman.afm
ZapfDingbats=ZapfDingbats.afm
.
FontPFM
Foobar-Bold=foobb__.pfm
% Example for an absolute path name with the prefix not applied (two equal signs)
Mistral==c:/psfonts/pfm/mist__.pfm
.
FontOutline
Code-128=Code_128.pfa
ArialMT=Arial.ttf
.
Encoding
iso8859-2=iso8859-2.enc
cp1250=cp1250.cpg
.
```

Searching for the UPR resource file. If only the built-in resources are to be used (PDF core fonts with host encoding), a UPR configuration file is not required, since PDFlib contains all necessary resources.

If other resources are to be used, PDFlib will search several places for a resource file. The process is configurable and consists of the following steps:

- ▶ The PDFlib ActiveX component checks a registry entry to find the file *pdflib.upr* in the *fonts* subdirectory of the PDFlib installation directory. If this file can't be opened, an *IOError* is raised.

- ▶ If a resource file can be opened during any of the above steps, but a required resource category cannot be found, a *SystemError* is raised.

Setting resources without a UPR file. In addition to using a UPR file for the configuration, it is also possible to directly configure individual resources within the source code via the `set_parameter()` function. This function 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:

```
oPDF.set_parameter "FontAFM", "Foobar-Bold=foobb__.afm"  
oPDF.set_parameter "FontOutline", "Foobar-Bold=foobb__.pfa"
```

Similar to UPR files, if two equal signs are present, the file name will be interpreted absolute. If only a single one equal sign is present, the directory prefix will be used if one has been configured.

3.3.7 CID Font Support for Japanese, Chinese, and Korean Text

CJK support in Acrobat and PDF¹. While Japanese font support was already available in Acrobat 3J, Acrobat 4 added full support for CID (Character ID) fonts for Japanese, Chinese, and Korean (CJK) text even in the non-Japanese versions of the full Acrobat package as well as the free Acrobat Reader. In order to use CJK documents in Acrobat you must do one of the following:

- ▶ Use a localized CJK version of Acrobat.
- ▶ If you use any non-CJK version of the full Acrobat product, select the Acrobat installer's option »Asian Language Support« (Windows) or »Language Kit« (Mac). The required support files (fonts and encodings) will be installed from the Acrobat product CD-ROM.
- ▶ If you use Acrobat Reader, install one of the Asian Font Packs which are available on the Acrobat 4 product CD-ROM, or on the Web.²

CJK encodings and fonts. Historically, a wide variety of CJK encoding schemes has been developed by diverse standards bodies, companies, and other organizations. Fortunately enough, all prevalent encodings are supported by Acrobat and PDF by default. Acrobat 4 supports a wealth of different encoding schemes for CJK fonts. Since the concept of an encoding is much more complicated for CJK text than for Latin text, simple encoding vectors with 256 entries no longer suffice. Instead, PostScript and PDF use the concept of character collections and character maps (CMaps) for organizing the characters in a font. Conceptually, CMaps can be thought of as large encodings for CJK fonts.

Acrobat 4 supports a set of standard fonts for CJK text. These fonts are supplied with the Acrobat installation (or the Asian FontPack), and therefore don't have to be embedded in the PDF file (this parallels the use of the 14 core fonts for Latin text). These fonts contain all characters required for common encodings, and support both horizontal and vertical writing modes. The standard fonts and CMaps are documented in Table 3.7. As can be seen from the table, the default CMaps support most CJK encodings used on Mac, Windows, and Unix systems, as well as several other vendor-specific encodings. In

1. This is a good opportunity to praise Ken Lunde's seminal tome »CJKV information processing – Chinese, Japanese, Korean & Vietnamese Computing« (O'Reilly 1999, ISBN 1-56592-224-7), as well as his work at Adobe since he's one of the driving forces behind CJK support in PostScript and PDF.

2. See <http://www.adobe.com/prodindex/acrobat/cjkfontpack.html>

particular, the major Japanese encoding schemes Shift-JIS, EUC, ISO 2022, and Unicode (UCS-2) are supported. Tables with all supported characters are available from Adobe¹; CMap descriptions can be found in Table 3.8.

Table 3.7. Acrobat’s standard fonts for Japanese, Chinese, and Korean text

locale	font name	font samples	supported CMaps (encodings)
Simplified Chinese	STSong-Light	国际	GB-EUC-H, GB-EUC-V, GBpc-EUC-H, GBpc-EUC-V, GBK-EUC-H, GBK-EUC-V, UniGB-UCS2-H, UniGB-UCS2-V
Traditional Chinese	MHei-Medium	中文	B5pc-H, B5pc-V, ETen-B5-H, ETen-B5-V, ETenms-B5-H, ETenms-B5-V, CNS-EUC-H, CNS-EUC-V, UniCNS-UCS2-H, UniCNS-UCS2-V
	MSung-Light	中文	
Japanese	HeiseiKakuGo-W5	日本語	83pv-RKSJ-H, 90ms-RKSJ-H, 90ms-RKSJ-V, 90msp-RKSJ-H, 90msp-RKSJ-V, 90pv-RKSJ-H, Add-RKSJ-H, Add-RKSJ-V, EUC-H, EUC-V, Ext-RKSJ-H, Ext-RKSJ-V, H, V, UniJIS-UCS2-H, UniJIS-UCS2-V, UniJIS-UCS2-HW-H, UniJIS-UCS2-HW-V
	HeiseiMin-W3	日本語	
Korean	HYGoThic-Medium	한국	KSC-EUC-H, KSC-EUC-V, KSCms-UHC-H, KSCms-UHC-V, KSCms-UHC-HW-H, KSCms-UHC-HW-V, KSCpc-EUC-H, UniKS-UCS2-H, UniKS-UCS2-V
	HYSMyeongJo-Medium	한국	

CJK font support in PDFlib. Having realized the similarity between core fonts/encoding vector on the one hand, and CJK standard fonts/CMaps on the other hand, it won’t be much of a surprise that both Latin and CJK fonts can be selected with the same PDFlib interface, using the CMap name in lieu of the encoding name, and taking into account that a given CJK font supports only a certain set of CMaps (see Table 3.7). For wide characters to work with the PDFlib ActiveX edition the *nativeunicode* parameter must be set to *true*. The *HeiseiKakuGo* sample in Table 3.7 has been generated with the following code:

```
oPDF.set_parameter "nativeunicode", "true"

font = oPDF.findfont("HeiseiKakuGo-W5", "Ext-RKSJ-H", 0)
oPDF.setfont font, 24
oPDF.set_text_pos x, y
oPDF.show ChrW(&H93FA) & ChrW(&H967B) & ChrW(&H8CEA)
```

These instructions locate one of the Japanese standard fonts, choosing a Shift-JIS-compatible CMap (*Ext-RKSJ*) encoding and horizontal writing mode (*H*). The *fontname* parameter must be the exact name of the font (strictly speaking, the value of the */CIDFont-Name* entry in the corresponding CID PostScript font file), without any encoding or writing mode suffixes. The *encoding* parameter is the name of one of the supported CMaps (the choice depends on the font) and will also indicate the writing mode (see below). PDFlib supports all of Acrobat’s default CMaps, and will complain when it detects a mismatch between the requested font and the CMap. For example, asking PDFlib to use a Korean font with a Japanese encoding will result in an exception of type *ValueError*.

Although CID font embedding is technically possible in PDF 1.3, it is not practical due to the size of typical CID fonts, and due to the fact that most CJK font licenses do not

1. See <http://partners.adobe.com/asn/developer/typeforum/cidfonts.html> for a wealth of resources related to CID fonts, including tables with all supported glyphs (search for «character collection»).

Table 3.8. Predefined CMaps for Japanese, Chinese, and Korean text (from the PDF Reference)

locale	supported CMaps	description	
Simplified Chinese	GB-EUC-H GB-EUC-V	Microsoft Code Page 936 (IfCharSet 0x86), GB 2312-80 character set, EUC-CN encoding	
	GBpc-EUC-H GBpc-EUC-V	Macintosh, GB 2312-80 character set, EUC-CN encoding, Script Manager code 2	
	GBK-EUC-H GBK-EUC-V	Microsoft Code Page 936 (IfCharSet 0x86), GBK character set, GBK encoding	
	UniGB-UCS2-H UniGB-UCS2-V	Unicode (UCS-2) encoding for the Adobe-GB1 character collection	
Traditional Chinese	B5pc-H B5pc-V	Macintosh, Big Five character set, Big Five encoding, Script Manager code 2	
	ETen-B5-H ETen-B5-V	Microsoft Code Page 950 (IfCharSet 0x88), Big Five character set with ETen extensions	
	ETenms-B5-H ETenms-B5-V	Same as ETen-B5-H, but replaces half-width Latin characters with proportional forms	
	CNS-EUC-H CNS-EUC-V	CNS 11643-1992 character set, EUC-TW encoding	
	UniCNS-UCS2-H UniCNS-UCS2-V	Unicode (UCS-2) encoding for the Adobe-CNS1 character collection	
	Japanese	83pv-RKSJ-H	Macintosh, JIS X 0208 character set with KanjiTalk6 extensions, Shift-JIS encoding, Script Manager code 1
90ms-RKSJ-H 90ms-RKSJ-V		Microsoft Code Page 932 (IfCharSet 0x80), JIS X 0208 character set with NEC and IBM extensions	
90msp-RKSJ-H 90msp-RKSJ-V		Same as 90ms-RKSJ-H, but replaces half-width Latin characters with proportional forms	
90pv-RKSJ-H		Macintosh, JIS X 0208 character set with KanjiTalk7 extensions, Shift-JIS encoding, Script Manager code 1	
Add-RKSJ-H Add-RKSJ-V		JIS X 0208 character set with Fujitsu FMR extensions, Shift-JIS encoding	
EUC-H EUC-V		JIS X 0208 character set, EUC-JP encoding	
Ext-RKSJ-H Ext-RKSJ-V		JIS C 6226 (JIS78) character set with NEC extensions, Shift-JIS encoding	
H V		JIS X 0208 character set, ISO-2022-JP encoding	
UniJIS-UCS2-H UniJIS-UCS2-V		Unicode (UCS-2) encoding for the Adobe-Japan1 character collection	
UniJIS-UCS2-HW-H UniJIS-UCS2-HW-V		Same as UniJIS-UCS2-H, but replaces proportional Latin characters with half-width forms	
Korean		KSC-EUC-H KSC-EUC-V	KS X 1001:1992 character set, EUC-KR encoding
		KSCms-UHC-H KSCms-UHC-V	Microsoft Code Page 949 (IfCharSet 0x81), KS X 1001:1992 character set plus 8822 additional hangul, Unified Hangul Code (UHC) encoding
	KSCms-UHC-HW-H KSCms-UHC-HW-V	Same as KSCms-UHC-H, but replaces proportional Latin characters with half-width forms	
	KSCpc-EUC-H	Macintosh, KS X 1001:1992 character set with Mac OS KH extensions, Script Manager Code 3	
	UniKS-UCS2-H UniKS-UCS2-V	Unicode (UCS-2) encoding for the Adobe-Korea1 character collection	

permit embedding. For this reason the *embed* parameter is not used for CID fonts, and must be 0.

PDFlib doesn't require any font-specific metrics information for CID fonts, and doesn't make any attempt to decode the client-supplied text strings, or verify whether they are correctly encoded with respect to the underlying CMap. For this reason the following features are currently not supported for CID fonts:

- ▶ calculating the extent of text with *stringwidth()*
- ▶ box formatting with *show_boxed()*
- ▶ activating underline/overline/strikeout mode
- ▶ retrieving the *textx/texty* position

Also, all characters in CJK fonts are considered to have the same width, including Latin characters. The character width is equal to the font size. If you want Latin characters which have a smaller width than the CJK characters you must switch to a Latin 8-bit font such as Courier or Helvetica.

Note PDFlib currently only supports the standard CID fonts supplied with Acrobat (see Table 3.7). Neither custom CID fonts nor Japanese, Chinese, or Korean TrueType fonts can be used. However, you can simulate bold fonts by rendering »filled and stroked« text (rendering mode 2, see text-rendering parameter).

Horizontal and vertical writing mode. PDFlib supports both horizontal and vertical writing modes. The mode is selected along with the encoding by choosing the appropriate CMap name. CMaps with names ending in *-H* select horizontal writing mode, while the *-V* suffix selects vertical writing mode.

Note Some PDFlib functions change their semantics according to the writing mode. For example, `continue_text()` should not be used in vertical writing mode, and the character spacing must be negative in order to spread characters apart in vertical writing mode. The details are discussed in the respective function descriptions.

CJK text encoding in PDFlib. The client is responsible for supplying text such that its encoding matches the encoding requested for the CID font. PDFlib does not check whether the supplied text conforms to the requested encoding. For multi-byte encodings, the high-order byte of a character must appear first.

PDFlib language bindings which are natively Unicode-aware (this includes ActiveX/COM) automatically convert Unicode strings supplied to the library. For this reason only Unicode-compatible CMaps should be used with these language bindings when the *nativeunicode* parameter is set to *true* (see also Section 3.3.8, »Unicode Support«).

Printing PDF documents with CJK text. Printing CJK documents gives rise to a number of issues which are outside the scope of this manual. However, we will supply some useful hints for the convenience of PDFlib users. If you have trouble printing CJK documents with Acrobat, consider one or more of the following:

- ▶ Printing CID fonts does not work on all PostScript printers. Native CID font support has only been integrated in PostScript version 2015, i.e. PostScript Level 1 and early Level 2 printers do not natively support CID fonts (unless the printer is equipped with the Type 0 font extensions). However, for early Level 2 devices the printer driver is supposed to take care of this by downloading an appropriate set of compatibility routines to pre-2015 Level 2 printers.

- ▶ Due to the large number of characters CID fonts consume very much printer memory (disk files for CID fonts typically are 5–10 MB in size). Not all printers have enough memory for printing such fonts. For example, in our testing we found that we had to upgrade a Level 3 laser printer from 16 MB to 48 MB RAM in order to reliably print PDF documents with CID fonts.
- ▶ Non-Japanese PostScript printers do not have any Japanese fonts installed. For this reason, you must check *Download Asian Fonts* in Acrobat's print dialog.
- ▶ If you can't successfully print using downloaded fonts, check *Print as Image* in Acrobat's print dialog. This instructs Acrobat to send a bitmapped version of the page to the printer (300 dpi, though).

3.3.8 Unicode Support

Starting with version 4, Acrobat supports the Unicode standard, almost identical to ISO 10646¹. This is a large character set which covers all current and many ancient languages and scripts in the world, and has significant support in many applications and operating systems. PDFlib supports the Unicode standard for the following features:



- ▶ bookmarks (see Figure 3.2)
- ▶ contents and title of note annotations (see Figure 3.2)
- ▶ standard and user-defined document information field contents (but not user-defined field names – the PDF specification unfortunately doesn't allow this)
- ▶ description and author of file attachments
- ▶ CJK text on page descriptions, provided a Unicode-compatible encoding is used (see Section 3.3.7, »CID Font Support for Japanese, Chinese, and Korean Text«)
- ▶ 8-bit code pages for TrueType and PostScript fonts

Before delving into the Unicode implementation, however, you should be aware of the following restrictions regarding Unicode support in Acrobat:

- ▶ Acrobat 4 does not display all characters from the Adobe Glyph List correctly (this has been fixed in Acrobat 5). This bug affects, for example, Cyrillic characters
- ▶ The usability of Unicode-enhanced PDF documents heavily depends on the Unicode support available on the target system. Unfortunately, most systems today are far from being fully Unicode-enabled in their default configurations. Although Windows NT and MacOS support Unicode internally, availability of appropriate Unicode fonts is still an issue.
- ▶ Acrobat on Windows is unable to handle more than one script in a single annotation. This seems to be related to an OS-specific issue (restrictions of the text edit widget used in Acrobat's implementation of the annotation feature).

Unicode code pages for PostScript and TrueType fonts. PDFlib supports Unicode for page descriptions for characters within the Adobe Glyph List (AGL). While text strings still must contain 8-bit characters, an arbitrary set of up to 256 characters can be selected using a Unicode-based code page definition file. This kind of Unicode support is available for Unicode-based TrueType fonts and PostScript with glyph names in the AGL. For details on code pages and AGL see Section 3.3.3, »Custom Encoding and Code Page Files for 8-Bit Encodings«.

¹ See <http://www.unicode.org> for more information about the Unicode standard

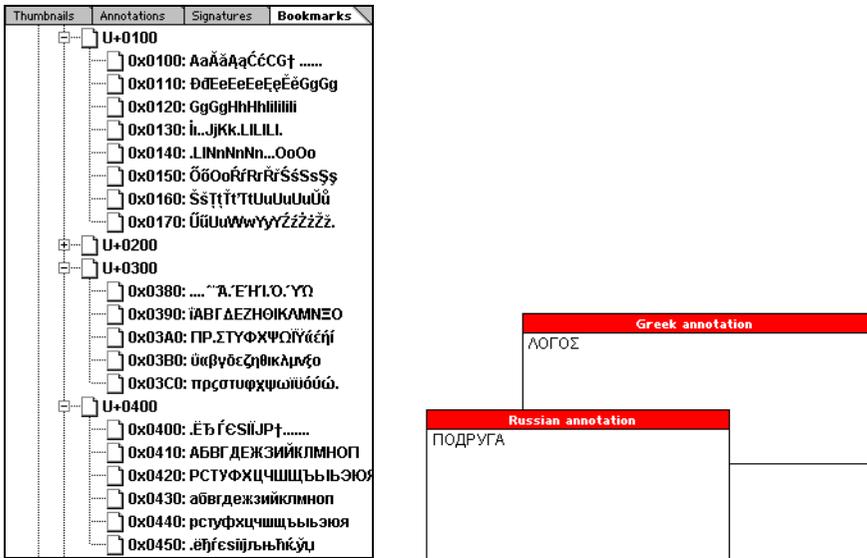


Fig. 3.2. Unicode bookmarks (left) and Unicode text annotations (right)

Unicode encoding for CID fonts. PDF allows Unicode-encoded text on document pages (as opposed to hypertext as discussed above). Unfortunately, this holds only true for CID fonts, but not regular Type 1 PostScript fonts. In order to place Unicode-conforming Chinese, Japanese, or Korean text on a page, a Unicode-compatible CMap must be used. These are easily identified by the *Uni* prefix in their name (see Table 3.8). These CMaps, however, only support the characters required for the respective locale, but not other Unicode characters.

Unicode encoding for hypertext elements. Users of the PDFlib ActiveX edition can directly supply Unicode text to the Unicode-enabled hypertext functions (bookmarks, annotations, etc.) when the *nativeunicode* parameter is set to *true*. For example, the following snippet (in hexadecimal notation) creates a bookmark with the Greek string »ΛΟΓΟΣ« (see Figure 3.2):

```
oPDF.set_parameter "nativeunicode", "true"
bookmark = oPDF.add_bookmark(ChrW(&H39B) & ChrW(&H39F) & ChrW(&H393) & ChrW(&H39F) &
ChrW(&H3A3), 0, 0)
```

Wrong Unicode character assignments on Windows. The following PDFlib language bindings are Unicode-aware, and can automatically convert Unicode strings to the format expected by PDFlib:

- ▶ ActiveX/COM
- ▶ Java
- ▶ Tcl (requires Tcl 8.2 or above)

However, in order to avoid the character conversion problem described below, Unicode support is disabled by default in these bindings. It can be activated by setting the PDFlib parameter *nativeunicode* to *true* (see also Section 4.3.2, »Text Output«):

```
oPDF.set_parameter "nativeunicode", "true"
```

Native Unicode mode means that the wrapper code will internally distinguish the following cases, and apply the appropriate conversion:

- ▶ 8-bit strings, i.e., strings which contain only characters from U+0000 to U+00FF are interpreted as PDFDocEncoding (for hypertext) or 8-bit characters according to the current encoding (for page descriptions).
- ▶ Unicode strings for hypertext functions will be encoded according to the PDF reference.
- ▶ Unicode strings for page descriptions will be supplied without any conversion. This requires a Unicode-compatible CMap to be selected (see Table 3.7).

The developer generally need not care about the encoding specifics detailed above, but can simply use Unicode text as supported by the environment. (More details on Unicode usage from within the supported languages can be found in the manual section for the respective binding in Chapter 2). However, there's a subtle issue related to literal Unicode characters embedded in ActiveX, Java, or Tcl source code which we will try to explain with a small example.

COM's native support for Unicode strings is just fine for PDF's hypertext elements, but can be dangerous with respect to page descriptions and non-Unicode-compliant 8-bit encodings. For example, while most characters in the Windows code page 1252 are compatible with Unicode, not all are (more specifically, the range &H80-&H9F). Consider the following attempt to show the endash character with PDFlib:

```
' Literal character &H96 = Alt-150 in the code. Works only if nativeunicode = false
oPDF.show("-");
```

If this snippet is used on a Windows system with code page 1252 and *nativeunicode* == *true*, the literal endash character (&H96 in code page 1252) will be translated to the corresponding Unicode character (&H2013 in this example), which is unsuited for an 8-bit PDF encoding such as *winansi*. In order to prevent this problem in native Unicode mode rewrite the above code snippet as follows:

```
' Safe way of selecting characters outside Latin-1 if nativeunicode = true
oPDF.show ChrW(&H96)
```

This will pass the intended character code &H96 to PDFlib, which will correctly interpret it according to the chosen encoding vector.

3.3.9 Text Metrics, Text Variations, and Text Box Formatting

Font and character metrics. PDFlib uses the character and font metrics system used by PostScript and 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 *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 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 values of capheight, ascender, and descender for a specific font are supplied in the font metrics file, and can be queried from PDFlib as follows:

```
font = oPDF.findfont("Times-Roman", "winansi", 0)
oPDF.setfont font, fontsize

capheight = oPDF.get_value("capheight", font) * fontsize
ascender = oPDF.get_value("ascender", font) * fontsize
descender = oPDF.get_value("descender", font) * fontsize
```

Note The position and size of superscript and subscript cannot be queried from PDFlib since this information is not contained in AFM metrics files.

CPI calculations. While most fonts have varying character widths, so-called mono-spaced fonts use the same widths for all characters. In order to relate PDF font metrics to the characters per inch (CPI) measurements often used in high-speed print environments, some calculation examples for the mono-spaced Courier font may be helpful. In Courier, all characters have a width of 600 units with respect to the full character cell of 1000 units per point (this value can be retrieved from the corresponding AFM metrics file). For example, with 12 point text all characters will have an absolute width of

$$12 \text{ points} * 600/1000 = 7.2 \text{ points}$$

with an optimal line spacing of 12 points. Since there are 72 points to an inch, exactly 10 characters of Courier 12 point will fit in an inch. In other words, 12 point Courier is a 10 cpi font. For 10 point text, the character width is 6 points, resulting in a $72/6 = 12$ cpi font. Similarly, 8 point Courier results in 15 cpi.

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. *set_parameter()* can be used to switch the underline, overline, and strikeout feature on or off as follows:

```
oPDF.set_parameter "underline", "true" ' enable underlines
```

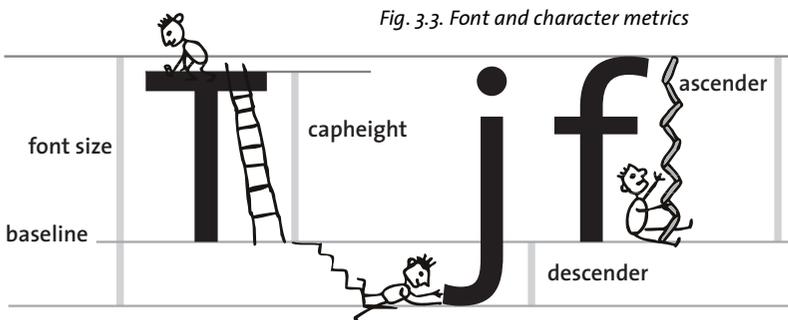


Fig. 3.3. Font and character metrics

The current stroke color is used for drawing the bars. The current linecap and dash parameters are ignored, however. Aesthetics alert: in most fonts underlining will touch descenders, and overlining will touch diacritical marks atop ascenders.

Note The underline, overline, and strikethrough features are not supported for CID fonts.

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 Table 3.9. They can be set with `set_value()`.

Table 3.9. Values for the text rendering mode

value	explanation	value	explanation
0	fill text	4	fill text and add it to the clipping path
1	stroke text (outline)	5	stroke text and add it to the clipping path
2	fill and stroke text	6	fill and stroke text and add it to the clipping path
3	invisible text	7	add text to the clipping path

```
oPDF.set_value "textrendering", 1 ' set stroked text rendering (outline text)
```

Text color. Text will usually be display in the current fill color, which can be set using `set_color()`. However, if a rendering mode other than 0 has been selected, both stroke and fill color may affect the text depending on the selected rendering mode.

Text box formatting. While PDFlib offers the `stringwidth()` function for performing text width calculations, many clients need easy access to text box formatting and justifying, e.g. to fit a certain amount of text into a given column. Although PDFlib offers such features, you shouldn't think of PDFlib as a full-featured text and graphics layout engine. The `show_boxed()` function is an easy-to-use method for text box formatting with a number of formatting options. Text may be laid out in a rectangular box either left-aligned, right-aligned, centered, or fully justified. The first line of text starts at a baseline with a vertical position which equals the top edge of the supplied box minus the leading. The bottom edge of the box serves as the last baseline used. For this reason, descenders of the last text line may appear outside the specified box (see Figure 3.4).

This function justifies by adjusting the inter-word spacing (the last line will be left-aligned only). Obviously, this requires that the text contains spaces (PDFlib will not insert spaces if the text doesn't contain any). Advanced text processing features such as hyphenation are not available – PDFlib simply breaks text lines at existing whitespace characters. Text is never clipped at the boundaries of the box.

Supplying a *feature* parameter of *blind* can be useful to determine whether a string fits in a given box, without actually producing any output.

ASCII newline characters (`ox0A`) in the supplied text are recognized, and force a new paragraph. CR/NL combinations are treated like a single newline character. Other formatting characters (especially tab characters) are not supported.

The following is a small example of using `show_boxed()`. It uses `rect()` to draw an additional border around the box which may be helpful in debugging:

In an attempt to reproduce sounds more accurately, pinyin spellings often differ markedly from the older ones, and personal names are usually spelled without apostrophes or hyphens; an apostrophe is sometimes used, however, to avoid ambiguity when syllables are run together (as in Chang'an to distinguish it from Chan'an).

Fig. 3.4. Top: Text box formatting: the bottom edge will serve as the last baseline, not as a clipping border. Right: text box formatting doesn't work if only a single word fits on a line. In the situation in the figure to the right, `show_boxed()` will not actually format any text.



```
Text = "In an attempt to reproduce sounds more accurately, pinyin spellings often ... "  
fontsize = 13
```

```
font = oPDF.findfont("Helvetica", "host", 0)  
oPDF.setfont font, fontsize
```

```
x = 50  
y = 650  
w = 357  
h = 6 * fontsize
```

```
c = oPDF.show_boxed(Text, x, y, w, h, "justify", "")  
If (c > 0) Then  
    ' Not all characters could be placed in the box; act appropriately here  
    ...  
End If  
oPDF.rect x, y, w, h  
oPDF.stroke
```

The following requirements and restrictions of `show_boxed()` shall be noted:

- ▶ Contiguous blanks in the text should be avoided.
- ▶ Due to restrictions in PDF's word spacing support, the *space* character must be available at code position &H20 in the encoding. Although this is the case for most common encodings, it implies that justification will not work with EBCDIC encoding.
- ▶ The simplistic formatting algorithm may fail for unsuitable combinations of long words and narrow columns. In particular, if only a single word fits in a column, `show_boxed()` will not format any text at all, but leave the column empty (see Figure 3.4).
- ▶ Since the bottom part of the box is used as a baseline, descenders in the last line may extend beyond the box area.
- ▶ Using `show_boxed()` with top-down coordinates isn't exactly intuitive. Please review the information in Section 3.2.1, »Coordinate Systems«.
- ▶ It's currently not possible to feed the text in multiple portions into the box formatting routine. However, you can retrieve the text position after calling `show_boxed()` with the `textx` and `texty` parameters.
- ▶ The font within the text box can't be changed.
- ▶ Text box formatting is not supported for CID fonts.

3.4 Image Handling

3.4.1 Supported Image File Formats

Embedding raster images in the generated PDF is an important feature of PDFlib. PDFlib currently deals with the image file formats described below. For most formats PDFlib passes the compressed image data unchanged to the PDF output since PDF internally supports most compression schemes used in image file formats. This technique (called *pass-through mode* in the descriptions below) results in very fast image import, since decompressing the image data and subsequent recompression are not necessary. However, PDFlib cannot check the integrity of the compressed image data in this mode. Incomplete or corrupt image data may result in error or warning messages when using the PDF document in Acrobat (e.g., »Read less image data than expected«).

If an image file can't be imported successfully and you need to know more details about the reason set the *imagewarning* parameter to *true* (see Section 4.6, »Image Functions« for more details):

```
oPDF.set_parameter "imagewarning", "true" ' enable image warnings
```

PNG images. PDFlib supports all flavors of PNG images (Portable Network Graphics).¹ PNG images are handled in pass-through mode in most cases. PNG images which make use of interlacing, contain an alpha channel (which will be lost anyway, see below), or have 16 bit color depth will have to be uncompressed, which takes significantly longer than pass-through mode. If a PNG image contains transparency information, the transparency is retained in the generated PDF (see Section 3.4.5, »Image Masks and Transparency«). Alpha channels are not supported by PDFlib.

JPEG images. JPEG images are always handled in pass-through mode. PDFlib supports the following flavors of JPEG images:

- ▶ The »baseline« JPEG flavor which accounts for the vast majority of JPEG images.
- ▶ Progressive JPEG compression which is supported since Acrobat 4/PDF 1.3. If run in Acrobat 3 compatibility mode PDFlib will refuse to import progressive JPEGs.

PDFlib applies a workaround which is necessary to correctly process Photoshop-generated CMYK JPEG files.

GIF images. GIF images are always handled in pass-through mode (PDFlib does not use LZW decompression). PDFlib supports the following flavors of GIF images:

- ▶ Due to restrictions in the compression schemes supported by the PDF file format, the entry in the GIF file called »LZW minimum code size« must have a value of 8 bits. Unfortunately, there is no easy way to determine this value for a certain GIF file. An image which contains more than 128 distinct color values will always qualify (e.g., a full 8-bit color palette with 256 entries). Images with a smaller number of distinct colors may also work, but it is difficult to tell in advance because graphics programs may use 8 bits or less as LZW minimum code size in this case, and PDFlib may therefore reject the image. The following trick which works in Adobe Photoshop and similar image processing software is known to result in GIF images which are accepted by PDFlib: load the GIF image, and change the image color mode from »indexed« to »RGB«. Now change the image color mode back to »indexed«, choosing a color pal-

1. See <http://www.w3.org/Graphics/PNG> and <http://www.libpng.org/pub/png>

ette with more than 128 entries, for example the Mac or Windows system palette, or the Web palette.

- ▶ The image must not be interlaced.
- ▶ Only the first image of a multi-frame (animated) GIF image will be imported.

For other GIF image flavors conversion to the PNG graphics format is recommended.

Note In a particular test case PDFlib converted a GIF image to a PDF file which displays just fine, but results in a PostScript error when printed to a PostScript Level 2 or 3 printer. Since the problem does not occur with Ghostscript, we consider this a bug in the PostScript interpreter. You can work around the problem by selecting PostScript Level 1 output in Acrobat's print dialog.

TIFF images. PDFlib will handle most TIFF images in pass-through mode. PDFlib supports the following flavors of TIFF images:

- ▶ compression schemes: uncompressed, CCITT (group 3, group 4, and RLE), ZIP (=Flate), LZW, and PackBits (=RunLength) are handled in pass-through mode; other compression schemes are handled by uncompressing.
- ▶ color depth: black and white, grayscale, RGB, and CMYK images; any alpha channel which may be present in the file is ignored.
- ▶ TIFF files containing more than one image (see Section 3.4.6, »Multi-Page Image Files«)
- ▶ Color depth must be 1, 2, 4, or 8 bits per color sample (this is a requirement of PDF)

Multi-strip TIFF images are converted to multiple images in the PDF file which will visually exactly represent the original image, but can be individually selected with Acrobat's image selection tool. Some TIFF features (e.g., CIE color space, JPEG compression) and certain combinations of features (e.g., LZW compression and alpha channel, LZW compression and tiling) are not supported.

Note Converting certain flavors of CCITT group 3 compressed TIFF images with PDFlib may trigger the message »Read less image data than expected« in Acrobat 4. Since the problem does not exist in Ghostscript or Acrobat 5, and the image displays just fine despite the error message, we consider this a bug in Acrobat 4. You may be able to work around it by choosing a different TIFF compression scheme.

CCITT images. Raw Group 3 or Group 4 fax compressed image data are always handled in pass-through mode. Note that this format actually means raw CCITT-compressed image data, *not* TIFF files using CCITT compression. Raw CCITT compressed image files are usually not supported in end-user applications, but can only be generated with fax-related software.

Raw data. Uncompressed (raw) image data may be useful for some special applications, e.g., constructing a color ramp directly in memory. The nature of the image is deduced from the number of color components: 1 component implies a grayscale image, 3 components an RGB image, and 4 components a CMYK image.

3.4.2 Code Fragments for Common Image Tasks

Embedding raster images with PDFlib is easy to accomplish. First, the image file has to be opened with a PDFlib function which does a brief analysis of the image parameters. The `open_image_file()` function returns a handle which serves as an image descriptor.

This handle can be used in a call to *place_image()*, along with positioning and scaling parameters:

```
lImage = oPDF.open_image_file("jpeg", "image.jpg", "", 0)
If (lImage = -1) Then
    MsgBox "Couldn't read image."
End
Else
    oPDF.place_image lImage, 0, 0, 1
    oPDF.close_image lImage
End If
```

The call to *close_image()* may or may not be required, depending on whether the same image will be used again in the same document (see Section 3.4.3, »Re-using Image Data«).

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 will be imported at 72 dpi if the user coordinate system hasn't been scaled (since there are 72 default units to an inch).

Resolution (dpi) values which may be contained in the original image file are ignored by PDFlib, but may be queried via the *resx* and *resy* parameters; the user is responsible for scaling the coordinate system appropriately (beware of non-square pixels). The following algorithm may be used to import an image at the resolution given in the file (or at 72 dpi if the image file doesn't contain any dpi value), and place it on the full page:

```
query the dpi values which may be present in the image file
dpi_x = oPDF.get_value("resx", lImage)
dpi_y = oPDF.get_value("resy", lImage)

' calculate scaling factors from the dpi values, see description of resx/resy
If (dpi_x > 0 And dpi_y > 0) Then      ' resx and resy are specified in the file
    scale_x = 72# / dpi_x
    scale_y = 72# / dpi_y
Else
If (dpi_x < 0 And dpi_y < 0) Then      ' only the ratio of resx and resy is known
    scale_x = 1#
    scale_y = dpi_y / dpi_x
Else                                    ' no information about resx and resy available
    scale_x = 1#
    scale_y = 1#
End If
End If

' create a new page such that the scaled image exactly fits, and place the image
oPDF.begin_page oPDF.get_value("imagewidth", lImage) * scale_x, _
                oPDF.get_value("imageheight", lImage) * scale_y
oPDF.[Scale] scale_x, scale_y
oPDF.place_image lImage, 0#, 0#, 1#
oPDF.close_image lImage
oPDF.end_page
```

In order to ignore any dpi value present in the image, and use a fixed dpi value instead (e.g. 300) replace the first two lines in the above code fragment with

```
dpi_x = 300
dpi_y = 300          ' or whatever you like
```

Forcing printed image size. In order to place an image on a PDF page such that it results in a specified target *width* and *height* (as opposed to specifying the resolution values as in the previous algorithm) with a lower left corner at (x, y) (all coordinates in points) the following algorithm may be used:

```
scale_x = lWidth / oPDF.get_value("imagewidth", lImage)
scale_y = lHeight / oPDF.get_value("imageheight", lImage)

oPDF.save

' scale the coordinate system to match the image size to the given rectangle
oPDF.[Scale] scale_x, scale_y

' in the positioning coordinates we must compensate for the above scaling
oPDF.place_image lImage, x / scale_x, y / scale_y, 1
oPDF.close_image lImage
oPDF.restore
```

Non-proportional image scaling. Since in most cases images will be scaled proportionally (i.e., using the same scaling factor in both dimensions), *place_image()* supports only a single scaling parameters which is applied to both dimensions. Non-proportional scaling can easily be achieved by scaling the coordinate system, bracketed with *save/restore* in order to not disturb other graphics operations. The following sequence will place an image, scaled to 50 percent horizontally and 75 percent vertically:

```
oPDF.save          ' save the original coordinate system
oPDF.[Scale] 0.5, 0.75 ' scale the coordinates, and therefore the image
oPDF.place_image lImage, 0#, 0#, 1#
oPDF.restore      ' restore the original coordinate system
```

Remember that the *x* and *y* positions supplied to *place_image()* will also be subject to the *scale()* call, and must be adjusted by dividing by the scaling factors.

A code fragment for placing images in a top-down coordinate system can be found in Section 3.2.1, »Coordinate Systems«.

3.4.3 Re-using Image Data

It should be emphasized that PDFlib supports an important PDF optimization technique for using repeated raster images.

Consider a layout with a constant logo or background on several pages. In this situation it is possible to include the image data only once in the PDF, and generate only a reference on each of the pages where the image is used. Simply open the image file and call *place_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.

3.4.4 Memory Images and External Image References

While the majority of image data for use with PDFlib will be pulled from some disk file on the local file system, other image data sources are also supported. For performance reasons supplying existing image data directly in memory may be preferable over opening a disk file. PDFlib supports in-core image data for certain image file formats.

PDFlib also supports an experimental feature which isn't recommended for general-use PDF files, but may offer advantages in certain environments. While almost all PDF documents are completely self-contained (the only exception being non-embedded fonts), it is also possible to store only a reference to some external data source in the PDF file instead of the actual image data, and rely on Acrobat to fetch the required image data when needed. This mechanism works similar to the well-known image references in HTML documents. Usable external image sources include data files in the local file system, and URLs. It is important to note that while file references work in Acrobat 3 and 4, URL references only work in Acrobat 4 or above (full product). PDF documents which include image URLs are neither usable in Acrobat 3 nor Acrobat Reader 4!

The `open_image()` interface can be used for both in-memory image data and external references.

3.4.5 Image Masks and Transparency

Transparency in PDF. Transparency has been missing from PostScript and PDF for quite a long time. Only with PDF 1.3 (and PostScript 3) Adobe integrated some limited support for transparency into languages and applications. While image masks (painting solid color through a bitmap mask) are an old feature of both PostScript and PDF, Acrobat 4 added the feature of masking particular pixels of an image. This offers the following opportunities:

- ▶ Masking by position: an image may carry the intrinsic information »print the foreground only, but not the background«. This is often used in catalog images.
- ▶ Masking by color value: pixels of a certain color (or from a color range – but not arbitrary sets of colors) are not painted, but the previously painted part of the page shines through instead. In TV and video technology this is also known as bluescreening, and is most often used for combining the weather man and the map into one image.

It is important to note that PDF supports binary transparency only: there is no alpha channel or variable opacity (»blend this image with the background«) but only a binary decision (»print either the image pixel, or the background pixel«). Binary transparency may be considered »poor man's alpha channel«. Another important restriction is that in PDF the mask is always attached to the image; it's not possible to use an image first with a mask, and the same image a second time without a mask, or with a different mask.

Viewing and printing PDF files with transparency. Equally important as PDF's intrinsic limitations with respect to transparency are the practical limitations when it comes to using PDF files with transparency in the viewer application. The following restrictions should be noted:

- ▶ Transparency only works in PDF 1.3/Acrobat 4 and above – older viewers will completely ignore transparency information, and display or print the whole image (overpainting the background).

- ▶ Printing transparent images to PostScript Level 1 or 2 doesn't work, even with Acrobat 4 (since transparency support only appeared in PostScript 3, and can't easily be emulated). Acrobat prints the base image without the mask.
- ▶ If an image is masked by position Acrobat 4 viewers will only honour the clipping up to a certain image size, and display the whole image otherwise. It appears from experimentation that the following limit applies to Acrobat 4 (Acrobat 5 is not affected by this limit):

$$\text{width} \times \text{height} \times \text{components} < 1024 \text{ K}$$

Images above this limit are displayed without applying the mask. The limit in a typical PostScript 3 printer seems to be lower, resulting in PostScript errors when trying to print PDF documents with large masked images.

Transparency support in PDFlib. PDFlib supports both masking by position and by color value (only single color values, but no ranges). Transparency information can be applied implicitly or explicitly. Masked images are not supported in Acrobat 3 compatibility mode.

In the implicit case, the transparency information from an external image file is respected, provided the image file format supports transparency or an alpha channel (this is not the case for all image file formats). Transparency information is detected in the following image file formats:

- ▶ GIF image files may contain a single transparent color value which is respected by PDFlib.
- ▶ PNG image files may contain several flavors of transparency information, or a full alpha channel. PDFlib tries to preserve as much as possible from this information: single transparent color values are retained; if multiple color values with an attached alpha value are given, only the first one with an alpha value below 50 percent is used; a full alpha channel is ignored.

The explicit case requires two steps, both of which involve image operations. First, an image must be prepared for later use as a binary transparency mask. This is accomplished by using the standard image file function with an additional parameter:

```
lMask = oPDF.open_image_file("png", MaskFileName, "mask", 0)
```

In order to be usable as a mask, an image must have only a single color component and a bit depth of 1, i.e., only plain bitmaps are suitable as a mask. Only PNG and in-memory images are supported for constructing a mask. Pixel values of 0 in the mask will result in the corresponding area of the image being painted, while pixel values of 1 result in the background shining through.

In the second step this mask is applied to another image which itself is acquired through one of the usual image functions:

```
lImage = oPDF.open_image_file(type, FileName, "masked", lMask)
If (lImage <> -1) Then
    oPDF.place_image lImage, 0#, 0#, 1#
Else
    ...
End If
```

Note the different use of the optional string parameter for `open_image_file()`: `mask` for defining a mask, and `masked` for applying a mask to another image. The integer parameter is unused in the first step, and carries the mask descriptor in the second step.

The image and the mask may have different pixel dimensions; the mask will automatically be scaled to the image's size.

PDFlib doesn't make any provisions for painting solid color through a mask (like PostScript's `imagemask` operator), since this is a special case of the general masking mechanism. You can achieve this effect by applying the required mask to an auxiliary image constructed in memory with `open_image()` (a solid rectangle of the requested color).

Note Multi-strip TIFF images are converted to multiple PDF images, which would be masked individually by PDFlib. Since this is usually not intended, this kind of images should be avoided as mask target. Also, it is important to not mix the implicit and explicit cases, i.e., don't use images with transparent color values as mask.

Ignoring transparency. Sometimes it is desirable to ignore any transparency information which may be contained in an image file. For example, Acrobat's anti-aliasing feature (also known as »smoothing«) isn't used for 1-bit images which contain black and transparent as their only colors. For this reason imported images with fine detail (e.g., rasterized text) may look ugly when the transparency information is retained in the generated PDF. In order to solve this problem, PDFlib's automatic transparency support can be disabled with the `ignoremask` parameter when opening the file:

```
lImage = oPDF.open_image_file("gif", FileName, "ignoremask", 0)
```

3.4.6 Multi-Page Image Files

PDFlib supports TIFF files which contain more than one image, also known as multi-page files. In order to use multi-page TIFFs, the call to `open_image_file()` additional string and numerical parameters are used:

```
lImage = oPDF.open_image_file("tiff", FileName, "page", 1)
```

The `page` parameter indicates that a multi-image file is to be used, and is only supported for TIFF images. The last parameter specifies the number of the image to use. The first image is numbered 1. This parameter may be increased until `open_image_file()` returns -1, signalling that no more images are available in the file.

A code fragment similar to the following can be used to convert all images in a multi-image TIFF file to a multi-page PDF file:

```
Frame = 1
Do
    lImage = oPDF.open_image_file("tiff", FileName, "page", Frame)
    If (lImage = -1) Then Exit Endif
    oPDF.begin_page lWidth, lHeight
    oPDF.place_image lImage, 0, 0, 1
    oPDF.close_image lImage
    oPDF.end_page
    Frame = Frame + 1
Loop
```

3.5 PDF Import with PDI

Note All functions described in this section require the additional PDF import library (PDI) which is not part of the PDFlib source code distribution. Although PDI is integrated into the ActiveX edition of PDFlib, PDFlib licensees must purchase an additional license key for PDI. Please visit our Web site for more information on obtaining PDI.

3.5.1 PDI Features and Applications

When the optional PDI (PDF import) library is attached to PDFlib, pages from existing PDF documents can be processed with all supported language bindings. The PDI product contains a parser for the PDF file format, and prepares pages from existing PDF documents for easy use with PDFlib. Conceptually, imported PDF pages are treated similarly to imported raster images such as TIFF or PNG: you open a PDF document, choose a page to import, and place it on an output page, applying any of PDFlib's transformation functions for translating, scaling, rotating, or skewing the imported page. Imported pages can easily be combined with new content by using any of 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:

- ▶ place a PDF background page and populate it with dynamic data (e.g., mail merge, personalized PDF documents on the Web, form filling)
- ▶ 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
- ▶ 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

3.5.2 Using PDI Functions with PDFlib

General considerations. It is important to understand that PDI will only import the actual page contents, but not any hypertext features (such as sound, movies, embedded files, hypertext links, form fields, bookmarks, thumbnails, and notes) which may be present in the imported PDF document. These hypertext features can be generated with the corresponding PDFlib functions. Similarly, you can not re-use individual elements 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 pages contain embedded font data for the same font, PDI will not remove any duplicate font data. On the other hand, if fonts are missing from some imported PDF, they will also be missing from the generated PDF output file.

PDFlib uses the template feature for placing imported PDF pages on the output page. Since some third-party PDF software does not correctly support the template feature,

restrictions in certain environments other than Acrobat may apply (see Section 3.2.4, »Templates«).

PDFlib-generated output which contains imported pages from other PDF documents can be processed with PDFlib/PDI again. However, due to restrictions in PostScript printing the nesting level should not exceed 10.

Code fragments 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):

```
Dim doc, page, pageno;
Dim sheetwidth, sheetheight;
Dim filename

...

pageno = 1

doc = oPDF.open_pdi(filename, "", 0)
If (doc = -1) Then
    MsgBox "Couldn't open input file!"
End
End If

page = oPDF.open_pdi_page(doc, pageno, "")
if (page = -1) Then
    MsgBox "Couldn't open page in input file!"
End
End If

sheetwidth = oPDF.get_pdi_value("width", doc, page, 0)
sheetheight = oPDF.get_pdi_value("height", doc, page, 0)

oPDF.begin_page sheetwidth, sheetheight
oPDF.place_pdi_page page, 0, 0, 1, 1
oPDF.close_pdi_page page

...add more content to the page using PDFlib functions...

oPDF.end_page
```

The PDFlib distribution contains PDI examples for all supported language bindings which demonstrate various applications of PDI features:

- ▶ The personalization demo pulls a page from an existing PDF document, and places additional text on the page.
- ▶ The quick reference demo extracts several pages from an existing PDF document, scales down the pages, and places multiple pages on an output sheet.
- ▶ The imposition demo (which is only available in C code) is a generalization of the quick reference demo. It processes an arbitrary number of PDF documents, and places $n \times m$ pages on an output sheet. In addition, lines are drawn around the scaled-down pages.

Dimensions of imported PDF pages. Imported PDF pages are regarded similarly to imported raster images, and can be placed on the output page using `place_pdi_page()`. PDI will import the page exactly as it is displayed in Acrobat, in particular:

- ▶ cropping will be retained (in technical terms: if a CropBox is present, PDI favors the CropBox over the MediaBox; see Section 3.2.2, »Page and Coordinate Limits«);
- ▶ rotation which has been applied to the page will be retained.

Many important properties, such as the page size of an imported PDF page and the number of pages in a document, can be queried via PDFlib's parameter mechanism. The relevant parameters are listed in Table 4.15 and Table 4.15. These properties can be useful in making decisions about the placement of imported PDF pages on the output page. The algorithms presented in Section 3.4.2, »Code Fragments for Common Image Tasks« for images can be used for scaling imported PDF pages as well.

Imported PDF pages and the graphics state. PDFlib treats imported PDF pages as templates. For this reason the comments in Section 3.2.4, »Templates« also apply to imported PDF pages which are placed on the output page. In particular, imported pages may change their appearance if the surrounding page changes some graphics state parameter which the imported page doesn't explicitly set. You can use `initgraphics()` to avoid this behavior.

Dealing with Acrobat 5/PDF 1.4 files. PDI is fully compatible to PDF 1.4 files generated with Acrobat 5. However, you should be aware of the following: Imported PDF documents must not have a higher PDF version number than the generated PDF output. Since the default output mode in PDFlib 4 is Acrobat 4/PDF 1.3, imported PDF 1.4 files will be rejected by default. You can modify this behavior by changing PDFlib's PDF output version as follows:

```
oPDF.set_parameter "compatibility", "1.4"
```

This will result in Acrobat 5/PDF 1.4 compatible output, which in turn allows you to also import files according to PDF 1.4.

3.5.3 Acceptable PDF Documents

Generally, PDI will happily process all kinds of PDF documents which can be opened with Acrobat, regardless of PDF version number or features used within the file. However, in rare cases a PDF document or a particular page of a document may be rejected by PDI. The following kinds of PDF documents can not be imported with PDI:

- ▶ PDF documents which use a higher PDF version number than the PDF output document that is currently being generated. The reason is that PDFlib can no longer make sure that the output will actually conform 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` parameter.
- ▶ Files with a damaged cross-reference table. You can identify such files by Acrobat's warning message *File is damaged but is being repaired*. Solution: open and resave the file with Acrobat.
- ▶ Encrypted PDF documents (i.e., any security settings applied). Solution: remove all security settings in Acrobat and resave the document. Obviously, you will need the document's password to do so.

- ▶ Since PDFlib/PDI do not contain any implementation of the LZW algorithm, certain PDF pages which use LZW compression (more specifically, LZW-compressed pages with multiple content streams) will be rejected (*unsupported filter*). Solution: resave the document in Acrobat 4.0 or above, with the *Optimize* button checked. Note that Acrobat 4.0 and above will never generate this kind of offended file, but only Acrobat 3 under certain circumstances. For this reason you are unlikely to run into this restriction. If you have a large number of such files which must be converted you should look into Acrobat's batch optimization feature.

Depending on the *pdiwarning* parameter, unacceptable PDF files will simply result in an error return value, or a nonfatal exception with a detailed explanation.

4 PDFlib API Reference

The API reference documents all supported PDFlib functions.

4.1 Data Types, Naming Conventions, and Scope

PDFlib data types. In general all ActiveX-aware development environments have access to PDFlib routines in the same way. In some cases you should be aware of different data types used for PDFlib function parameters. In particular, the data types used for strings and binary data may differ slightly. Table 4.1 lists useful information about these data types.

Table 4.1. Data types in the language bindings

development environment	string data type	binary data type
ActiveX in general	BSTR (string)	variant of type VT_ARRAY VT_Uh ¹
Delphi	String (for 8-bit encodings) or WideString (for Unicode)	OleVariant

1. In other words, a variant array of unsigned bytes.

Function scopes. Most PDFlib functions are subject to certain ordering and nesting constraints which are derived from their contribution to the generated document. Most of these constraints are rather obvious. For example, you must begin a page before you can close it. In the same spirit, the functions for opening a PDF document and closing it must always be paired. PDFlib uses a strict scoping system for defining and verifying the correct ordering of functions used by client programs. The function descriptions reference these scopes; the scope definitions can be found in Table 4.2. Figure 4.1 depicts the relationship of scopes. PDFlib will throw an exception if a function is called outside the allowed scope.

Table 4.2. Function scope definitions

scope name	definition
path	started by one of <code>moveto()</code> , <code>circle()</code> , <code>arc()</code> , <code>arcn()</code> , or <code>rect()</code> terminated by any of the functions in Section 4.4.4, »Path Painting and Clipping«
page	between <code>begin_page()</code> and <code>end_page()</code> , but outside of path scope
template	between <code>begin_template()</code> and <code>end_template()</code> , but outside of path scope
pattern	between <code>begin_pattern()</code> and <code>end_pattern()</code> , but outside of path scope
document	between <code>open_file()</code> and <code>close()</code> , but outside of page, template, and pattern scope
object	scopeAnytime during the lifetime of the PDFlib object, but outside of document

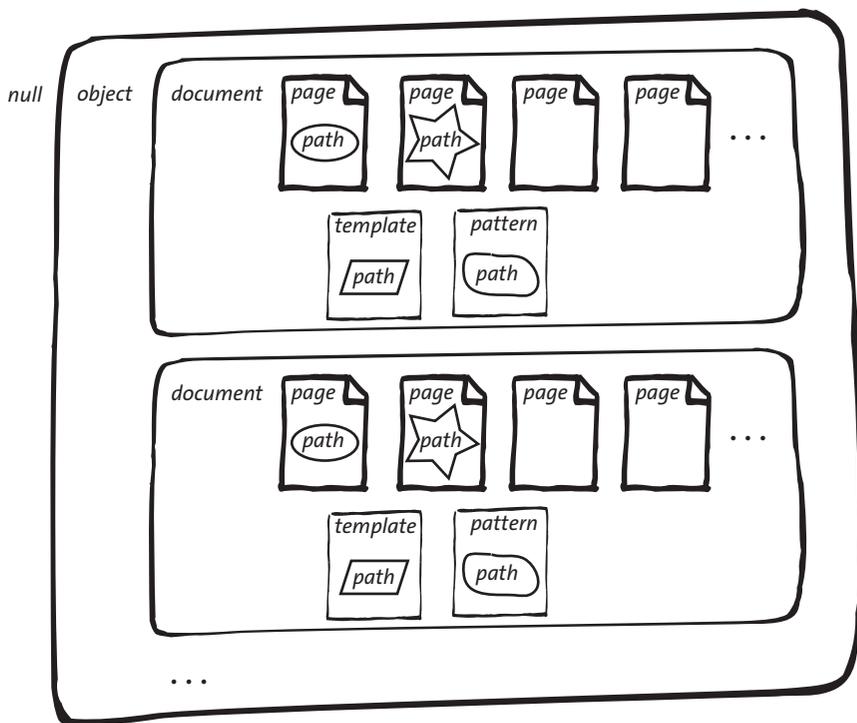


Fig. 4.1.
Relationship of scopes

4.2 General Functions

4.2.1 Setup

Table 4.3 lists relevant parameters and values for this section.

Table 4.3. Parameters and values for the setup functions

function	key	explanation
set_parameter	compatibility	Set PDFlib's compatibility mode to one of the strings »1.2«, »1.3«, or »1.4« for Acrobat 3, 4, or 5. The default is »1.3«. This parameter must be set before the first call to open_*(<i>i</i>). Setting compatibility to »1.2« will make Acrobat 4 features unavailable. Strict Acrobat 3 compatibility mode is not required for generating Acrobat 3 compatible files, but only under very specific circumstances related to PDF-enabled RIPs (see Section 1.3, »PDFlib Output and Compatibility«).
set_parameter	prefix	Resource file name prefix as used in a UPR file (see Section 3.3.6, »Resource Configuration and the UPR Resource File«). The prefix can only be set once. It contains a slash character plus a path name, which in turn may start with a slash.
set_parameter	resourcefile	Relative or absolute file name of the PDFlib UPR resource file. The resource file will be loaded at the next attempt to access resources. The resource file name can only be set once. This call should occur before the first page.
set_parameter	serial	Set the PDFlib and/or PDI serial string (see Section 2.2.2, »Installing the PDFlib ActiveX Edition«). The serial string can only be set once before the first call to begin_page(<i>i</i>)
set_parameter	warning	Enable or suppress warnings (nonfatal exceptions). Possible values are true and false, default value is true.

Table 4.3. Parameters and values for the setup functions

function	key	explanation	
set_value	compress	Set the compression parameter to a value from 0–9. This parameter does not affect precompressed image data which is handled in pass-through mode.	
		0	no compression
		1	best speed
		6	default value
		9	best compression
get_value	major	Return the major, minor, or revision number of PDFlib, respectively.	
	minor		
	revision		
get_parameter	version	Return the full PDFlib version string in the format <major>.<minor>.<revision>, possibly suffixed with additional qualifiers such as beta, rc, etc.	

4.2.2 Document and Page

Table 4.4 lists relevant parameters and values for this section.

Table 4.4. Parameters and values for the document and page functions

function	key	explanation
set_value	pagewidth	Change the page size dimensions of the current page. These parameters must only be used within a page description.
	pageheight	
set_value	CropBox,	Change one of the box parameters of the current page. These parameters must only be used within a page description. The parameter name must be followed by a slash '/' character and one of llx, lly, urx, ury, for example: CropBox/lx (see Section 3.2.2, »Page and Coordinate Limits« for details)
	BleedBox,	
	ArtBox,	
	TrimBox	

Function open_file(filename As String) As Long

Create a new PDF file using the supplied file name.

filename Name of the PDF output file to be generated. If *filename* is empty the PDF document will be generated in memory instead of on file. The result must be fetched by the client with the *get_buffer()* function. *open_file()* will always succeed in this case, and never return the -1 error value.

Returns -1 on error, and 1 otherwise.

Details This function creates a new PDF file using the supplied *filename*. PDFlib will attempt to open a file with the given name, and close the file when the PDF document is finished.

Scope *object*; this function starts document scope if the file could successfully be opened, and must always be paired with a matching *close()* call.

Function get_buffer()

Get the contents of the PDF output buffer. The result must be used by the client before calling any other PDFlib function.

Returns A buffer full of binary PDF data for consumption by the client. The returned buffer can be used until the end of the surrounding *object* scope.

Details Fetch the full or partial buffer containing the generated PDF data. If this function is called between page descriptions, it will return the PDF data generated so far. If it is called after *close()* it returns the remainder of the PDF document. If there is only a single call to this function which happens after *close()* the returned buffer is guaranteed to contain the complete PDF document in a contiguous buffer.

Since PDF output contains binary characters, client software must be prepared to accept non-printable characters including null values.

Scope *object, document* (in other words: after *end_page()* and before *begin_page()*, or after *close()* and before *delete()*). This function can only be used if an empty filename has been supplied to *open_file()*.

Sub *close()*

Close the generated PDF file, and release all document-related resources.

Details This function finishes the generated PDF document, free all document-related resources, and close the output file if the PDF document has been opened with *open_file()*. This function must be called when the client is done generating pages, regardless of the method used to open the PDF document.

When the document was generated in memory (as opposed to on file), the document buffer will still be kept after this function is called (so that it can be fetched with *get_buffer()*), and will be freed in the next call to *open()*, or when the PDFlib object goes out of scope.

Scope *document*; this function terminates document scope, and must always be paired with a matching call *open()* functions.

Sub *begin_page(width As Single, height As Single)*

Add a new page to the document.

width, height The *width* and *height* parameters are the dimensions of the new page in points. Acrobat's page size limits are documented in Section 3.2.1, »Coordinate Systems«. A list of commonly used page formats can be found in Table 4.24 in Section 4.9, »Page Size Formats«. The page size can be changed after calling *begin_page()* with the *pagewidth* and *pageheight* parameters. In order to produce landscape pages use *width > height*. PDFlib uses *width* and *height* to construct the page's MediaBox. You can use several parameters to set other box entries in the PDF (see Table 4.3).

Scope *document*; this function starts page scope, and must always be paired with a matching *end_page()* call.

Params *pagewidth, pageheight, CropBox, BleedBox, ArtBox, TrimBox*

Sub *end_page()*

Finish the page.

Details This function must be used to finish a page description.

Scope *page*; this function terminates page scope, and must always be paired with a matching *begin_page()* call.

4.2.3 Parameter Handling

PDFlib maintains a number of internal parameters which are used for controlling PDFlib's operation and the appearance of the PDF output. Four functions are available for setting and retrieving both numerical and string parameters. All parameters (both keys and values) are case-sensitive. The descriptions of available parameters can be found in the respective sections.

Function `get_value(key As String, modifier As Single) As Single`

Get the value of some PDFlib parameter with numerical type.

key The name of the parameter to be queried.

modifier An optional modifier to be applied to the parameter. Whether a modifier is required and what it relates to is explained in the various parameter tables. If the modifier is unused it must be 0.

Returns The numerical value of the parameter.

Scope Depends on *key*.

See also `get_pdi_value()`

Sub `set_value(key As String, value As Single)`

Set the value of some PDFlib parameter with numerical type.

key The name of the parameter to be set.

value The new value of the parameter to be set.

Scope Depends on *key*.

Function `get_parameter(key As String, modifier As Single) As String`

Get the contents of some PDFlib parameter with string type.

key The name of the parameter to be queried.

modifier An optional modifier to be applied to the parameter. Whether a modifier is required and what it relates to is explained in the various parameter tables. If the modifier is unused it must be 0.

Returns The string value of the parameter. The returned string can be used until the end of the surrounding *document* scope.

Scope Depends on *key*.

See also `get_pdi_parameter()`

Sub set_parameter(key As String, value As String)

Set some PDFlib parameter with string type.

key The name of the parameter to be set.

value The new value of the parameter to be set.

Scope Depends on *key*.

4.3 Text Functions

4.3.1 Font Handling

Table 4.5 lists relevant parameters and values for this section.

Table 4.5. Parameters and values for the font functions (see Section 4.2.3, «Parameter Handling»)

function	key	explanation
<i>set_parameter</i>	<i>FontAFM</i> <i>FontPFM</i> <i>FontOutline</i> <i>Encoding</i>	<i>The corresponding resource file line as it would appear for the respective category in a UPR file (see Section 3.3.6, «Resource Configuration and the UPR Resource File»). Multiple calls add new entries to the internal list. (See also prefix and resourcefile in Table 4.3)</i>
<i>get_value</i>	<i>font</i>	<i>Return the identifier of the current font which must have been set with setfont(). Scope: page, pattern, template.</i>
<i>get_parameter</i>	<i>fontname</i>	<i>The name of the current font which must have been previously set with setfont(). Scope: page, pattern, template.</i>
<i>get_parameter</i>	<i>fontencoding</i>	<i>The name of the encoding or CMap used with the current font. A font must have been previously set with setfont(). Scope: page, pattern, template.</i>
<i>get_value</i>	<i>fontsize</i>	<i>Return the size of the current font which must have been previously set with setfont(). Scope: page, pattern, template.</i>
<i>get_value</i>	<i>capheight</i> <i>ascender</i> <i>descender</i>	<i>Return metrics information for the font identified by the modifier. See Section 3.3.9, «Text Metrics, Text Variations, and Text Box Formatting» for more details. The values are measured in fractions of the font size, and must therefore be multiplied by the desired font size.</i>
<i>set_parameter</i>	<i>fontwarning</i>	<i>If set to false, findfont() returns -1 if the font/encoding combination cannot be loaded (instead of throwing an exception). Default is true.</i>

Function findfont(fontname As String, encoding As String, embed As Long) As Long

Search for a font, and prepare it for later use.

fontname The name of the font as configured in PDFlib.

encoding For 8-bit fonts, *encoding* is one of *builtin*, *macroman*, *winansi*, *ebcdic*, or *host* (see Section 3.3.2, «8-Bit Encodings built into PDFlib»), or the name of an external PDFlib-supplied or user-defined encoding (see Section 3.3.3, «Custom Encoding and Code Page Files for 8-Bit Encodings»). Note that in order to use arbitrary encodings, you will need metrics information for the font (see Section 3.3.5, «PostScript and TrueType Fonts»).

Alternatively, *encoding* can be the name of one of the built-in CMaps if *fontname* describes a CID font (see Section 3.3.7, «CID Font Support for Japanese, Chinese, and Korean

Text»). In this case metrics information is not required. Case is significant for both *fontname* and *encoding*.

embed If 0 (zero), only general font information (a font descriptor) is included in the PDF output. If 1, the font outline file must be available in addition to the metrics information, and the actual font definition will be embedded in the PDF output. However, the font file will only be checked when this function is called, but not yet used, since font embedding is done at the end of the generated PDF file. The *embed* parameter must be 0 for CID fonts.

Returns A font handle for later use with *setfont()*. The behavior of this function changes when the *fontwarning* parameter is set to *true*. In this case *findfont()* returns an error code of -1 if the requested font/encoding combination cannot be loaded, and does not throw an exception. However, exceptions will still be thrown when bad parameters are passed.

The returned number – the font handle – doesn't have any significance to the user other than serving as an argument to *setfont()* and related functions. In particular, requesting the same font/encoding combination in different documents may result in different font handles.

Details This function prepares a font for later use with *setfont()*. The metrics will be loaded from memory or from an external metrics file. If the requested font/encoding combination cannot be used due to configuration problem (e.g., a font, metrics, or encoding file could not be found, or a mismatch was detected), an exception of type *RuntimeError* will be raised. Otherwise, the value returned by this function can be used as font argument to other font-related functions.

CID fonts are not supported in Acrobat 3 compatibility mode.

Scope *document, page, pattern, template*

Params See Table 4.5.

Sub *setfont(font As Long, fontsize As Single)*

Set the current font in the given size.

font A font handle returned by *findfont()*.

fontsize Size of the font, measured in units of the current user coordinate system.

Details The font must be set on each page before drawing any text. Font settings will not be retained across pages. The current font can be changed an arbitrary number of times per page.

Scope *page, pattern, template*

Params See Table 4.5.

4.3.2 Text Output

Note All text supplied to the functions in this section must match the encoding selected with *findfont()*. This applies to 8-bit text as well as Unicode or other encodings selected via a CMap.

Table 4.5 lists relevant parameters and values for this section. The scope for all parameters is *page, pattern, and template* unless otherwise noted.

Table 4.6. Parameters and values for the text functions (see Section 4.2.3, »Parameter Handling«)

function	key	explanation
set_value get_value	leading	Set or get the leading, which is the distance between baselines of adjacent lines of text. The leading is used for <code>continue_text()</code> and set to the value of the font size when a new font is selected using <code>setfont()</code> . Setting the leading equal to the font size results in dense line spacing. However, ascenders and descenders of adjacent lines will generally not overlap.
set_value get_value	textrise	Set or get the text rise parameter. The text rise specifies the distance between the desired text position and the default baseline. Positive values of text rise move the baseline up. The text rise always relates to the vertical coordinate. This may be useful for superscripts and subscripts. The text rise is set to the default value of 0 at the beginning of each page.
set_value get_value	horizscaling	Set or get the horizontal text scaling to the given percentage, which must be greater than 0. Text scaling shrinks or expands the text by a given percentage. The text scaling is set to the default of 100 at the beginning of each page. Text scaling always relates to the horizontal coordinate.
set_value get_value	textrendering	Set or get the current text rendering mode to one of the values given in Table 3.9. The text rendering parameter is set to the default of 0 (= solid fill) at the beginning of each page.
set_value get_value	charspacing	Set or get the character spacing, i.e., the shift of the current point after placing the individual characters in a string. The spacing is given in text space units. It is reset to the default of 0 at the beginning of each page. In order to spread the characters apart use positive values for horizontal writing mode, and negative values for vertical writing mode.
set_value get_value	wordspacing	Set or get the word spacing, i.e., the shift of the current point after placing individual words in a text line. In other words, the current point is moved horizontally after each ASCII space character (&H20). Since fonts with multi-byte encodings don't have an ASCII space character they are not affected by the word spacing. The spacing value is given in text space units. It is reset to the default of 0 at the beginning of each page.
get_value	textx texty	Get the x or y coordinate, respectively, of the current text position. These parameters are currently not supported for CID fonts.
set_parameter get_parameter	underline overline strikeout	Set or get the current underline, overline, and strikeout modes, which are retained until they are explicitly changed. These modes can be set independently from each other, and are reset to false at the beginning of each page (see Section 3.3.9, »Text Metrics, Text Variations, and Text Box Formatting«). true underline/overline/strikeout text (does not work for CID fonts) false do not underline/overline/strikeout text
set_parameter	native- unicode	If true, enable native Unicode text processing for language bindings with Unicode support; disable if false. Default value is false (see Section 3.3.8, »Unicode Support«). Scope: any.

Sub show(text As String)

Print text in the current font and size at the current position.

text The text to be printed.

Details Both font (via `setfont()`) and current text position (via `set_text_pos()` or some text output function) must have been set before. The current point is moved to the end of the printed text.

Scope page, pattern, template

Params See Table 4.6.

Sub show_xy(text As String, x As Single, y As Single)

Print *text* in the current font at position (*x*, *y*).

text The text to be printed.

x,y The position in the user coordinate system where the text will be printed.

Details The font must have been set before with *setFont()*. The current point is moved to the end of the printed text.

Scope *page, pattern, template*

Params See Table 4.6.

Sub continue_text(text As String)

Print text at the next line.

text The text to be printed.

Details The positioning of text and the spacing between lines is determined by the *leading* parameter and the most recent call to *show_xy()* or *set_text_pos()*. This function can also be used after *show_boxed()* if that function has been called with *mode = left* or *justify*. The current point is moved to the end of the printed text.

Scope *page, pattern, template*; this function should not be used in vertical writing mode.

Params See Table 4.6.

Function show_boxed(text As String, x As Single, y As Single, width As Single, height As Single, mode As String, feature As String) As Long

Format text into a text box according to the requested formatting mode.

text The text to be formatted into the box.

x, y The coordinates of a corner of the text box or the coordinates of the alignment point if *width = 0* and *height = 0*.

width, height The size of the text box, or 0 for single-line formatting.

mode *mode* selects the horizontal alignment mode. If *width = 0* and *height = 0*, *mode* can attain one of the values *left*, *right*, or *center*, and the text will be formatted according to the chosen alignment with respect to the point (*x*, *y*), with *y* denoting the position of the baseline. In this mode, this function does not check whether the submitted parameters result in some text being clipped at the page edges, nor does it apply any line-wrapping. It returns the value 0 in this case.

If *width* or *height* is different from 0, *mode* can attain one of the values *left*, *right*, *center*, *justify*, or *fulljustify*. The supplied text will be formatted into a text box defined by the lower left corner (*x*, *y*) (but see the description of top-down coordinates in Section 3.2.1, »Coordinate Systems«) and the supplied *width* and *height*. If the text doesn't fit into

a line, a simple line-breaking algorithm is used to break the text into the next available line, using existing space characters for possible line-breaks. While the *left*, *right*, and *center* modes align the text on the respective line, *justify* aligns the text on both left and right margins. According to common practice the very last line in the box will only be left-aligned in *justify* mode, while in *fulljustify* mode all lines (including the last one if it contains at least one space character) will be left- and right-aligned. *fulljustify* is useful if the text is to be continued in another column.

feature If the *feature* parameter is *blind*, all calculations are performed (with the exception of the internal *textx* and *texty* coordinates, which are not updated), but no text output is actually generated. This can be used for size calculations and possibly trying different font sizes for fitting some amount of text into a given box by varying the font size. Otherwise *feature* must be empty.

Returns The number of characters which could not be processed since the text didn't completely fit into the column. If the text did actually fit, it returns 0. Since no formatting is performed if *width* = 0 and *height* = 0, the function always returns 0 in this case.

Details The current font must have been set before calling this function. The current values of font, font size, horizontal spacing, and leading are used for the text.

Scope *page, pattern, template*; this function cannot be used with CID fonts or *ebcdic* encoding.

Params See Table 4.6.

See also Restrictions of this functions are listed in Section 3.3.9, »Text Metrics, Text Variations, and Text Box Formatting«.

Function `stringwidth(text As String, font As Long, size As Single) As Single`

Return the width of *text* in an arbitrary font.

text The text for which the width will be queried.

font A font handle returned by *findfont()*.

size Text size, measured in units of the user coordinate system.

Details This function returns the width of *text* in an arbitrary font and size which has been selected with *findfont()*. The width calculation takes the current values of the following text parameters into account: horizontal scaling, character spacing, and word spacing.

Scope *page, pattern, template, path, document*; this function cannot be used with CID fonts. If the current font is a CID font, it returns 0 regardless of the *text* and *size* arguments.

Params See Table 4.6.

Sub `set_text_pos(x As Single, y As Single)`

Set the text output position.

x, y The current text position to be set.

Details The text position is set to the default value of (o, o) at the beginning of each page. The current point for graphics output and the current text position are maintained separately.

Scope *page, pattern, template*

Params See Table 4.6.

4.4 Graphics Functions

4.4.1 General Graphics State

Note None of the the general graphics state functions must be used during path scope (see Section 3.2, »Page Descriptions«).

Sub **setdash**(*b As Single, w As Single*)

Set the current dash pattern.

b, w The number of alternating black and white units. *b* and *w* must be non-negative numbers.

Details In order to produce a solid line, set $b = w = o$. The dash parameter is set to solid at the beginning of each page.

Scope *page, pattern, template*

Sub **setpolydash**(*darray*)

Set a more complicated dash pattern defined by an array.

darray An array which contains alternating values for black and white dash lengths. The array values must be non-negative, and not all zero.

Details In order to produce a solid line, choose an empty array. The array length must be less than or equal to 8; otherwise the array will be truncated. The dash parameter is set to a solid line at the beginning of each page.

Scope *page, pattern, template*

Sub **setflat**(*flatness As Single*)

Set the flatness parameter.

flatness Describes the maximum distance (in device pixels) between the path and an approximation constructed from straight line segments.

Details The flatness parameter is set to the default value of o at the beginning of each page, which means that the device's default flatness is used.

Scope *page, pattern, template*

Sub `setlinejoin(linejoin As Long)`

Set the `linejoin` parameter.

linejoin Specifies the shape at the corners of paths that are stroked, see Table 4.7.

Details The `linejoin` parameter is set to the default value of 0 at the beginning of each page.

Scope `page, pattern, template`

Table 4.7. Values of the `linejoin` parameter

value	description (from the PDF reference)	examples
0	Miter joins: the outer edges of the strokes for the two segments are continued until they meet. If the extension projects too far, as determined by the miter limit, a bevel join is used instead.	
1	Round joins: a circular arc with a diameter equal to the line width is drawn around the point where the segments meet and filled in, producing a rounded corner.	
2	Bevel joins: the two path segments are drawn with butt end caps (see the discussion of <code>linecap</code> parameter), and the resulting notch beyond the ends of the segments is filled in with a triangle.	

Sub `setlinecap(linecap As Long)`

Set the `linecap` parameter.

linecap Controls the shape at the end of a path with respect to stroking, see Table 4.8.

Details The `linecap` parameter is set to the default value of 0 at the beginning of each page.

Scope `page, pattern, template`

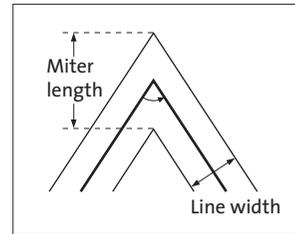
Table 4.8. Values of the `linecap` parameter

value	description (from the PDF reference)	examples
0	Butt end caps: the stroke is squared off at the endpoint of the path.	
1	Round end caps: a semicircular arc with a diameter equal to the line width is drawn around the endpoint and filled in.	
2	Projecting square end caps: the stroke extends beyond the end of the line by a distance which is half the line width and is squared off.	

Sub `setmiterlimit(miter As Single)`

Set the miter limit.

miter A value greater than or equal to 1.



Details If the `linejoin` parameter is set to 0 (miter join), two line segments joining at a small angle will result in a sharp spike. This spike will be replaced by a straight end (i.e., the miter join will be changed to a bevel join) when the ratio of the miter length and the line width exceeds the miter limit. The miter limit is set to the default value of 10 at the beginning of each page. This corresponds to an angle of roughly 11.5 degrees.

Scope *page, pattern, template*

Sub `setlinewidth(width As Single)`

Set the current line width.

width The line width in units of the current user coordinate system.

Details The `width` parameter is set to the default value of 1 at the beginning of each page.

Scope *page, pattern, template*

Sub `initgraphics`

Reset all color and graphics state parameters to their defaults.

Details The color, linewidth, linecap, linejoin, miterlimit, dash parameter, and the current transformation matrix (but not the text state parameters) are reset to their respective defaults.

This function may be useful in situations where the program flow doesn't allow for easy use of `save()/restore()`, or for preparing the graphics state for a subsequent template or imported PDF.

Scope *page, pattern, template*

4.4.2 Special Graphics State

All graphics state parameters are restored to their default values at the beginning of a page. The default values are documented in the respective function descriptions. Functions related to the text state are listed in Section 4.3, »Text Functions«.

All transformation functions (`translate()`, `scale()`, `rotate()`, `skew()`, `concat()`, `setmatrix()`, and `initgraphics()`) change the coordinate system used for drawing subsequent objects. They do not affect existing objects on the page at all.

Sub save()

Save the current graphics state.

Details The graphics state contains parameters that control all types of graphics objects. Saving the graphics state is not required by PDF; it is only necessary if the application wishes to return to some specific graphics state later (e.g., a custom coordinate system) without setting all relevant parameters explicitly again. The following items are subject to *save/restore*:

- ▶ graphics parameters: clipping path, coordinate system, current point, flatness, line cap style, dash pattern, line join style, line width, miter limit;
- ▶ color parameters: fill and stroke colors;
- ▶ some PDFlib parameters, see list below.

Pairs of *save()* and *restore()* may be nested. Although the PDF specification doesn't limit the nesting level of *save/restore* pairs, applications must keep the nesting level below 10 in order to avoid printing problems caused by restrictions in the PostScript output produced by PDF viewers, and to allow for additional save levels required by PDFlib internally.

Scope *page, pattern, template*; must always be paired with a matching *restore()* call. *save()* and *restore()* calls must be balanced on each page, pattern, and template.

Params The following parameters are subject to *save/restore*: *charspacing, wordspacing, horizscaling, leading, font, fontsize, textrendering, textrise*;

The following parameters are not subject to *save/restore*: *fillrule, underline, overline, strikeout*.

Sub restore()

Restore the most recently saved graphics state.

Details The corresponding graphics state must have been saved on the same page, pattern, or template.

Scope *page, pattern, template*; must always be paired with a matching *save()* call. *save()* and *restore()* calls must be balanced on each page, pattern, and template.

Sub translate(tx As Single, ty As Single)

Translate the origin of the coordinate system.

tx, ty The new origin of the coordinate system is the point *(tx, ty)*, measured in the old coordinate system.

Scope *page, pattern, template*

Sub scale(sx As Single, sy As Single)

Scale the coordinate system.

sx, sy Scaling factors in x and y direction.

Details This function scales the coordinate system by *sx* and *sy*. It may also be used for achieving a reflection (mirroring) by using a negative scaling factor. One unit in the *x* direction in the new coordinate system equals *sx* units in the *x* direction in the old coordinate system; analogous for *y* coordinates.

Visual Basic users please observe Section 2.2.7, »Using PDFlib with Visual Basic«, and use the notation *oPDF.[scale]*.

Scope *page, pattern, template*

Sub rotate(phi As Single)

Rotate the user coordinate system.

phi The rotation angle in degrees.

Details Angles are measured counterclockwise from the positive *x* axis of the current coordinate system. The new coordinate axes result from rotating the old coordinate axes by *phi* degrees.

Scope *page, pattern, template*

Sub skew(alpha As Single, beta As Single)

Skew the coordinate system.

alpha, beta Skewing angles in *x* and *y* direction in degrees.

Details Skewing (or shearing) distorts the coordinate system by the given angles in *x* and *y* direction. *alpha* is measured counterclockwise from the positive *x* axis of the current coordinate system, *beta* is measured clockwise from the positive *y* axes. Both angles must be in the range $360^\circ < \alpha, \beta < 360^\circ$, and must be different from $-270^\circ, -90^\circ, 90^\circ$, and 270° .

Scope *page, pattern, template*

Sub concat(a As Single, b As Single, c As Single, d As Single, e As Single, f As Single)

Concatenate a matrix to the current transformation matrix.

a, b, c, d, e, f Elements of a transformation matrix. The six floating point values make up a matrix in the same way as in PostScript and PDF (see references). In order to avoid degenerate transformations, $a*d$ must not be equal to $b*c$.

Details This function concatenates a matrix to the current transformation matrix (CTM) for text and graphics. It allows for the most general form of transformations. Unless you are familiar with the use of transformation matrices, the use of *translate()*, *scale()*, *rotate()*, and *skew()* is suggested instead of this function. The CTM is reset to the identity matrix $[1, 0, 0, 1, 0, 0]$ at the beginning of each page.

Scope *page, pattern, template*

Sub setmatrix(a As Single, b As Single, c As Single, d As Single, e As Single, f As Single)

Explicitly set the current transformation matrix.

a, b, c, d, e, f See *concat()*.

Details This function is similar to *concat()*. However, it disposes of the current transformation matrix, and completely replaces it with a new matrix.

Scope *page, pattern, template*

4.4.3 Path Construction

Table 4.9 lists relevant parameters and values for this section.

Table 4.9. Parameters and values for the path segment functions (see Section 4.2.3, »Parameter Handling«)

function	key	explanation
<i>get_value</i>	<i>currentx</i> <i>currenty</i>	The x or y coordinate, respectively, of the current point.

Note Make sure to call one of the functions in Section 4.4.4, »Path Painting and Clipping« after using the functions in this section, or the constructed path will have no effect, and subsequent operations may raise a PDFlib exception.

Sub moveto(x As Single, y As Single)

Set the current point.

x, y The coordinates of the new current point.

Details The current point is set to the default value of *undefined* at the beginning of each page. The current points for graphics and the current text position are maintained separately.

Scope *page, pattern, template, path*; this function starts *path* scope.

Params *currentx, currenty*

Sub lineto(x As Single, y As Single)

Draw a line from the current point to another point.

x, y The coordinates of the second endpoint of the line.

Details This function adds a straight line from the current point to (x, y) to the current path. The current point must be set before using this function. The point (x, y) becomes the new current point.

The line will be centered around the »ideal« line, i.e. half of the linewidth (as determined by the value of the linewidth parameter) will be painted on each side of the line connecting both endpoints. The behavior at the endpoints is determined by the value of the linecap parameter.

Scope *path*

Params *currentx, currenty*

Sub curveto(x1 As Single, y1 As Single, x2 As Single, y2 As Single, x3 As Single, y3 As Single)

Draw a Bézier curve from the current point, using three more control points.

x1, y1, x2, y2, x3, y3 The coordinates of three control points.

Details A Bézier curve is added to the current path from the current point to (x_3, y_3) , using (x_1, y_1) and (x_2, y_2) as control points. The current point must be set before using this function. The endpoint of the curve becomes the new current point.

Scope *path*

Params *currentx, currenty*

Sub circle(x As Single, y As Single, r As Single)

Draw a circle.

x, y The coordinates of the center of the circle.

r The radius of the circle.

Details This function adds a circle to the current path as a complete subpath. The point $(x + r, y)$ becomes the new current point. The resulting shape will be circular in user coordinates. If the coordinate system has been scaled differently in x and y directions, the resulting curve will be elliptical.

Visual Basic users please observe Section 2.2.7, »Using PDFlib with Visual Basic«, and use the notation *oPDF[**circle**]*.

Scope *page, pattern, template, path*; this function starts *path* scope.

Params *currentx, currenty*

Sub arc(x As Single, y As Single, r As Single, alpha As Single, beta As Single)

Draw a counterclockwise circular arc segment.

x, y The coordinates of the center of the circular arc segment.

r The radius of the circular arc segment. *r* must be nonnegative.

alpha, beta The start and end angles of the circular arc segment in degrees.

Details This function adds a counterclockwise circular arc segment to the current path, extending from *alpha* to *beta* degrees. For both functions, angles are measured counterclockwise from the positive x axis of the current coordinate system. If there is a current point an additional straight line is drawn from the current point to the starting point of the arc. The endpoint of the arc becomes the new current point.

The arc segment will be circular in user coordinates. If the coordinate system has been scaled differently in x and y directions the resulting curve will be elliptical.

Scope *page, pattern, template, path*; this function starts *path* scope.

Params *currentx, currenty*

Sub arc(x As Single, y As Single, r As Single, alpha As Single, beta As Single)

Like *PDF_arc()*, but draws a clockwise circular arc segment.

Details Except for the drawing direction, this function behave exactly like *PDF_arc()*. In particular, the angles are still measured *counterclockwise* from the positive x axis.

Sub rect(x As Single, y As Single, width As Single, height As Single)

Draw a rectangle.

x, y The coordinates of the lower left corner of the rectangle.

width, height The size of the rectangle.

Details This function adds a rectangle to the current path as a complete subpath. Setting the current point is not required before using this function. The point (x, y) becomes the new current point. The lines will be centered around the »ideal« line, i.e. half of the linewidth (as determined by the value of the linewidth parameter) will be painted on each side of the line connecting the respective endpoints.

Scope *page, pattern, template, path*; this function starts *path* scope.

Params *currentx, currenty*

Sub closepath()

Close the current path.

Details This function closes the current subpath, i.e., adds a line from the current point to the starting point of the subpath.

Scope *path*

Params *currentx, currenty*

4.4.4 Path Painting and Clipping

Table 4.10 lists relevant parameters and values for this section.

Table 4.10. Parameters and values for the path painting and clipping functions (see Section 4.2.3, »Parameter Handling«)

function	key	explanation
<i>set_parameter</i>	<i>fillrule</i>	Set the current fill rule to <i>winding</i> or <i>evenodd</i> . The fill rule is used by PDF viewers to determine the interior of shapes for the purpose of filling or clipping. Since both algorithms yield the same result for simple shapes, most applications won't have to change the fill rule. The fill rule is reset to the default of <i>winding</i> at the beginning of each page.

Note All functions in this section clear the path, and leave the current point undefined. Subsequent drawing operations must explicitly set the current point (e.g., using *moveto()*) after one of these functions has been called.

Sub stroke()

Stroke the path and clear it.

Details This function strokes (draws) the current path with the current line width and the current stroke color.

Scope *path*; this function terminates *path* scope.

Sub closepath_stroke()

Close the path, and stroke it.

Details This function closes the current subpath (adds a straight line segment from the current point to the starting point of the path), and strokes the complete current path with the current line width and the current stroke color.

Scope *path*; this function terminates *path* scope.

Sub fill()

Fill the interior of the path with the current fill color.

Details This function fills the interior of the current path with the current fill color. The interior of the path is determined by one of two algorithms (see *setfillrule()*). Open paths are implicitly closed before being filled.

Scope *path*; this function terminates *path* scope.

Params *fillrule*

Sub fill_stroke()

Fill and stroke the path with the current fill and stroke color.

Details This function fills and strokes the current path with the current fill and stroke color, respectively.

Scope *path*; this function terminates *path* scope.

Params *fillrule*

Sub closepath_fill_stroke()

Close the path, fill, and stroke it.

Details This function closes the current subpath (adds a straight line segment from the current point to the starting point of the path), and fills and strokes the complete current path.

Scope *path*; this function terminates *path* scope.

Params *fillrule*

Sub clip()

Use the current path as clipping path.

Details This function uses the intersection of the current path and the current clipping path as the clipping path for subsequent operations. The clipping path is set to the default value of the page size at the beginning of each page. The clipping path is subject to *save()/restore()*. It can only be enlarged by means of *save()/restore()*.

Scope *path*; this function terminates *path* scope.

Sub endpath()

End the current path without filling or stroking it.

Details This function doesn't have any visible effect on the page and will only rarely be useful. It generates an invisible path on the page.

Scope *path*; this function terminates *path* scope.

4.5 Color Functions

Sub setcolor(*type* As String, *colorspace* as String, *c1* As Single, *c2* As Single, *c3* As Single, *c4* As Single)

Set the current color space and color.

type One of *stroke*, *fill*, or *both* to specify that the color is set for filling, stroking, or both filling and stroking.

colorspace One of *gray*, *rgb*, *cmyp*, *spot*, or *pattern* to specify the color space.

c1, c2, c3, c4 Color components for the chosen color space:

- ▶ If the *colorspace* is *gray*, *c1* specifies a gray value;
- ▶ If the *colorspace* is *rgb*, *c1*, *c2*, *c3* specify red, green, and blue values;
- ▶ If the *colorspace* is *cmyp*, *c1*, *c2*, *c3*, *c4* specify cyan, magenta, yellow, and black values;
- ▶ If *colorspace* is *spot*, *c1* specifies a spot color handle returned by *makespotcolor()*, and *c2* specifies a tint value between 0 and 1;
- ▶ If *colorspace* is *pattern*, *c1* specifies a pattern handle returned by *begin_pattern()*.

Details All color values for the *gray*, *rgb*, and *cmyp* color spaces and the tint value for the *spot* color space must be numbers in the inclusive range 0–1. Unused parameters should be set to 0.

Grayscale, RGB values and spot color tints are interpreted according to additive color mixture, i.e., 0 means no color and 1 means full intensity. Therefore, a gray value of 0 and RGB values with $(r, g, b) = (0, 0, 0)$ mean black; a gray value of 1 and RGB values with $(r, g, b) = (1, 1, 1)$ mean white. CMYK values, however, are interpreted according to subtractive color mixture, i.e., $(c, m, y, k) = (0, 0, 0, 0)$ means white and $(c, m, y, k) = (0, 0, 0, 1)$ means black. Color values in the range 0–255 must be scaled to the range 0–1 by dividing by 255.

The fill and stroke color values for the *gray*, *rgb*, and *cmymk* color spaces are set to a default value of black at the beginning of each page. There are no defaults for spot and pattern colors.

Scope *page, pattern (only if the pattern's paint type is 1), template*; a pattern color can not be used within its own definition.

Function `makespotcolor(spotname As String) As Long`

Make a named spot color from the current fill color.

spotname An arbitrary name for the spot color to be defined. This name may contain arbitrary characters, but is restricted to a maximum length of 126 bytes.

Returns A color handle which can be used in subsequent calls to `setcolor()` throughout the document. Spot color handles can be reused across all pages, but not across documents. There is no limit for the number of spot colors in a document.

Details If *spotname* has already been used in a previous call to `makespotcolor*()`, the return value will be the same as in the earlier call, and will not reflect the current color.

The special spot color name *All* can be used to apply color to all color separations, which is useful for painting registration marks. A spot color name of *None* will produce no visible output on any color separation.

Scope *page, pattern, template*; the current fill color must not be a spot color or pattern;

Function `begin_pattern(width As Single, height As Single, xstep As Single, ystep As Single, painttype As Long) As Long`

Start a pattern definition.

width, height The dimensions of the pattern's bounding box in points.

xstep, ystep The offsets when repeatedly placing the pattern to stroke or fill some object. Most applications will set these to the pattern *width* and *height*, respectively.

painttype If *painttype* is 1 the pattern must contain its own color specification which will be applied when the pattern is used; if *painttype* is 2 the pattern must not contain any color specification but instead the current fill or stroke color will be applied when the pattern is used for filling or stroking.

Returns A pattern handle that can be used in subsequent calls to `setcolor()` throughout the document.

Details Hypertext functions and functions for opening images must not be used during a pattern definition, but all text, graphics, and color functions (with the exception of the pattern which is in the process of being defined) can be used.

Scope *document*; this function starts *pattern* scope, and must always be paired with a matching `end_pattern()` call.

Sub end_pattern(pattern As Long)

Finish a pattern definition.

Scope *pattern*; this function terminates *pattern* scope, and must always be paired with a matching *begin_pattern()* call.

Sub setgray_fill(g As Single)**Sub setgray_stroke(g As Single)****Sub setgray(g As Single)**

Deprecated, use *setcolor(type, "gray", g, o, o, o)* with a type parameter of *fill*, *stroke*, or *both* instead.

Sub setrgbcolor_fill(red As Single, green As Single, blue As Single)**Sub setrgbcolor_stroke(red As Single, green As Single, blue As Single)****Sub setrgbcolor(red As Single, green As Single, blue As Single)**

Deprecated, use *setcolor(type, "rgb", red, green, blue, o)* with a type parameter of *fill*, *stroke*, or *both* instead.

4.6 Image Functions

The functions for opening images described below can be called within or outside of page descriptions. Opening images outside a *begin_page()* / *end_page()* context actually offers slight output size advantages.

Table 4.11 lists relevant parameters and values for this section.

Table 4.11. Parameters and values for the image functions (see Section 4.2.3, »Parameter Handling«)

function	key	explanation
<i>get_value</i>	<i>imagewidth</i> <i>imageheight</i>	Get the width or height, respectively, of an image in pixels. The modifier is the integer handle of the selected image.
<i>get_value</i>	<i>resx</i> <i>resy</i>	Get the horizontal or vertical resolution of an image, respectively. The modifier is the integer handle of the selected image. If the value is positive, the return value is the image resolution in pixels per inch (dpi). If the return value is negative, it can be used to find the aspect ratio of non-square pixels, but doesn't have any absolute meaning. If the return value is zero, the resolution of the image is unknown.
<i>set_parameter</i>	<i>image-warning</i>	This parameter can be used in order to obtain more detailed information about why an image couldn't be opened successfully with <i>open_image_file()</i> or <i>open_CCITT()</i> : <i>true</i> Raise a Nonfatal exception when the image function fails. The message supplied with the exception may be useful in debugging. <i>false</i> Do not raise an exception when the image function fails. Instead, the function returns -1 on error. This is the default.

Function open_image_file(type As String, filename As String, stringparam As String, intparam As Long) As Long

Open an image file.

type Specifies the format type of the image: *png*, *gif*, *jpeg*, or *tiff* (see Section 3.4.1, »Supported Image File Formats«). Case is significant for all parameters.

filename The name of the image file to be opened.

stringparam, intparam The *stringparam* and *intparam* parameters are used for additional image attributes according to Table 4.12. If *stringparam* is unused, it must be an empty string, and *intparam* must be 0.

Returns An image handle which can be used in subsequent image-related calls. The return value must be checked for -1 which signals an error. In order to get more detailed information about the nature of an image-related problem (wrong image file name, unsupported format, bad image data, etc.), set the *imagewarning* parameter to *true* (see Table 4.11). The returned image handle can not be reused across multiple PDF documents.

Details This function opens and analyzes a raster graphics file in one of the supported file formats as determined by the *type* parameter. PDFlib will open the image file with the given name, process the contents, and close it before returning from this call. Although images can be placed multiply within a document (see *place_image()*), the actual image file will not be kept open after this call.

Scope *document, page*; must always be paired with a matching *close_image()* call.

Params *imagewidth, imageheight, resx, resy, imagewarning*

Table 4.12. The *stringparam* and *intparam* parameters of *open_image_file()*

stringparam	explanation and possible intparam values
<i>mask</i>	Create a mask from this image. The returned image handle may be used in subsequent calls for opening another image and can be supplied for the »masked« parameter. The <i>intparam</i> parameter is ignored in this case, and must be 0.
<i>masked</i>	Use the image descriptor given in <i>intparam</i> as a mask for this image. The <i>intparam</i> parameter is an image handle which has been retrieved with a previous call to <i>open_image()</i> with the »mask« parameter, and has not yet been closed.
<i>ignoremask</i>	Ignore any transparency information which may be present in the image file.
<i>invert</i>	Invert black and white for 1-bit TIFF images. This is mainly intended as a workaround for certain TIFF images which are interpreted differently by different applications.
<i>page</i>	Extract the image with the number given in <i>intparam</i> from a multi-page file. The first image has the number 1. This is only supported for multi-image TIFF files.

Function open_CCITT(filename As String, width As Long, height As Long, BitReverse As Long, K As Long, BlackIs1 As Long) As Long

Open a raw CCITT image.

filename The name of the CCITT file to be opened.

width, height The dimensions of the image in pixels.

BitReverse If 1, do a bitwise reversal of all bytes in the compressed data.

K CCITT compression parameter for encoding scheme selection.

- ▶ -1 indicates G4 encoding;
- ▶ 0 indicates one-dimensional G3 encoding (G3-1D);
- ▶ 1 indicates mixed one- and two-dimensional encoding (G3, 2-D) as supported by PDF.

BlackIs1 If this parameter has the value 1, 1-bits are interpreted as black and 0-bits as white. Most CCITT images don't use such a black-and-white reversal, i.e., most images use *BlackIs1* = 0.

Returns An image handle which may be used in subsequent image-related calls if not -1. Since PDFlib is unable to analyze CCITT images, all relevant image parameters have to be passed to *open_CCITT()* by the client.

Details This function opens an image file with raw CCITT G3 or G4 compressed bitmap data (this is different from a TIFF file which contains CCITT-compressed image data!).

Scope *document, page*; must always be paired with a matching *close_image()* call.

Params *imagewidth, imageheight, resx, resy, imagewarning*

Function *open_image*(type As String, source As String, data, length As Long, width As Long, height As Long, components As Long, bpc As Long, params As String) As Long

Use image data from a variety of data sources.

type Specifies the kind of image data or compression: *jpeg, ccitt, or raw* (see Section 3.4.1, »Supported Image File Formats«).

source, data, length The *source* parameter denotes where the image data comes from, and can attain the values *fileref, url, or memory* (see Section 3.4.4, »Memory Images and External Image References«). The relationship among the *source, data, and length* parameters is explained in Table 4.13.

width, height The dimensions of the image in pixels.

components The number of color components must be 1, 3, or 4 corresponding to grayscale, RGB, or CMYK image data.

bpc The number of bits per component must be 1, 2, 4, or 8.

params If *components* = 1 and *bpc* = 1, *params* may be *mask* in order to use this image as an image mask. Alternatively, additional CCITT parameters can be supplied (see below).

Returns An image handle which may be used in subsequent image-related calls if not -1

Details This versatile interface can be used to work with image data in several formats and from several data sources. Unlike *open_image_file()* which analyzes an image file, the user must supply the *length, width, height, components, and bpc* parameters. *open_image()* does not analyze the image data, and the user is responsible for supplying parameters which actually match the image properties. Otherwise corrupt PDF output may be generated, and Acrobat may respond with the *Image in Form, Type 3 font or pattern too big* error message.

Table 4.13. Values of the source, data, and length parameters of `open_image()`

source	data	length
<code>fileref^f</code>	string ² with a platform-independent file name (see [1])	unused, should be 0
<code>url^l</code>	string ² with an image URL conforming to RFC 1738 ³ . The URL will not be resolved by PDFlib, but by Acrobat when the PDF is opened (see Section 3.4.4, »Memory Images and External Image References«). This experimental feature is not recommended for production use.	unused, should be 0
<code>memory</code>	Binary bytes containing image data; the image data is compressed according to the <code>type</code> parameter. Exactly »length« bytes must be supplied.	length of (compressed) image data in bytes.

1. Not supported in Acrobat 3 compatibility mode.
2. data is not a string but an array of bytes, which makes it a little bit clumsy to pass filenames or URLs.
3. The URL must not contain any additional parameter, query string, access scheme, network login, or fragment identifier.

If `type` is `raw`, length must be equal to $[width \times components \times bpc / 8] \times height$ bytes, with the bracketed term adjusted upwards to the next integer, and this exact amount of data must be supplied. The image samples are expected in the standard PostScript/PDF ordering, i.e., top to bottom and left to right (assuming no coordinate transformations have been applied). Even if `bpc` is not 8, each pixel row begins on a byte boundary, and color values must be packed from left to right within a byte. Image samples are always interleaved, i.e., all color values for the first pixel are supplied first, followed by all color values for the second pixel, and so on.

If `type` is `ccitt`, CCITT-compressed image data is expected. In this case, `params` is examined. For CCITT images two parameters as described for `open_CCITT()` can be supplied in the `params` string as follows:

```
/K -1 /BlackIs1 true
```

Supported values for `/K` are -1, 0, or 1, the default value is 0. Supported values for `/BlackIs1` are `true` and `false`; the default value is `false`. The default values will be used if an empty `params` string is supplied. `BitReverse` cannot be supplied in this string. Instead, a special notion is used: if `length` is negative, the image data will be reversed.

If `params` is not used, it must be empty. The client is responsible for the memory pointed to by the `data` argument. The memory may be freed by the client immediately after this call.

Don't use Photoshop-generated CMYK JPEG images with this function since these will appear in the PDF with inverted colors.

Scope `document, page`; must always be paired with a matching `close_image()` call.

Params `imagewidth, imageheight, resx, resy, imagewarning`

Sub `close_image(image As Long)`

Close an image.

image A valid image handle retrieved with one of the `open_image*()` functions.

Details This function only affects PDFlib's associated internal image structure. If the image has been opened from file, the actual image file is not affected by this call since it has already been closed at the end of the corresponding `open_image*()` call. An image handle cannot be used any more after it has been closed with this function, since it breaks PD-

Flib's internal association with the image. This function must not be called for image handles which have been opened with *open_pdi_page()* (use *PDF_close_pdi_page()* instead).

Scope *document, page*; must always be paired with a matching call to one of the *open_image_file()*, *open_CCITT()*, *open_image()* functions.

Sub place_image(image As Long, x As Single, y As Single, scale As Single)

Place an image or template, with the lower left corner at (x, y) , and scale it.

image A valid image handle retrieved with one of the *open_image*()* or *begin_template()* functions.

x, y The coordinates in the user coordinate system where the lower left corner of the placed image will be located.

scale The scaling factor which will be applied to the image in *x* and *y* direction.

Details See Section 3.4.2, »Code Fragments for Common Image Tasks« for more information on scaling and dpi calculations, including non-uniform scaling (different scaling factors in *x* and *y* dimensions).

Scope *page, pattern, template*; this function can be called an arbitrary number of times on arbitrary pages, as long as the image handle has not been closed with *close_image()*.

See also *initgraphics()* may be useful for initializing the graphics state before placing templates.

Function begin_template(width As Long, height As Long) As Long

Start a template definition.

width, height The dimensions of the template's bounding box in points.

Returns An image handle which can be used in subsequent image-related calls. There is no error return.

Details Hypertext functions and functions for opening images must not be used during a template definition, but all text, graphics, and color functions can be used.

Scope *document*; this function starts *template* scope, and must always be paired with a matching *end_template()* call.

Params *imagewidth, imageheight*

Sub end_template()

Finish a template definition.

Scope *template*; this function terminates *template* scope, and must always be paired with a matching *begin_template()* call.

4.7 PDF Import Functions

4.7.1 Document and Page

Function `open_pdi(filename As String, stringparam As String, intparam As Long) As Long`

Open an existing PDF document and prepare it for later use.

filename The name of the PDF file.

stringparam, intparam Reserved for later use; must currently be or empty and 0, respectively.

Returns A document descriptor which can be used for processing individual pages of the document or for querying document properties. A return value of -1 indicates that the PDF document couldn't be opened. An arbitrary number of PDF documents can be opened simultaneously. The return value can be used until the end of the enclosing document scope.

Details In order to get more detailed information about the nature of a PDF-related problem (wrong PDF file name, unsupported format, bad PDF data, etc.), set the *pdiwarning* parameter to *true*.

Scope *document, page*

Params See Table 4.15 and Table 4.14.

Sub `close_pdi(doc As Long)`

Close all open PDI page handles, and close the input PDF document.

doc A valid PDF document handle retrieved with *open_pdi()*.

Details This function closes a PDF import document, and releases all resources related to the document. All document pages which may be open are implicitly closed. The document handle must not be used after this call.

Scope *document, page*

Params See Table 4.14 and Table 4.15.

Function `open_pdi_page(doc As Long, pagenumber As Long, pagelabel As String) As Long`

Prepare a page for later use with *place_pdi_page()*.

doc A valid PDF document handle retrieved with *open_pdi()*.

pagenumber The number of the page to be opened. The first page has page number 1.

pagelabel Reserved; must currently be NULL or empty.

Returns A page descriptor which can be used for placing pages with *place_image()*. A return value of -1 indicates that the page couldn't be opened. The return value can be used until the end of the enclosing document scope.

Details In order to get more detailed information about a problem related to PDF import (unsupported format, bad PDF data, etc.), set the *pdwarning* parameter to *true*.

An arbitrary number of pages can be opened simultaneously. If the same page is opened multiply, different handles will be returned, and each handle must be closed exactly once. Opening the same page more than once is not recommended because the actual page data will be copied to the output document more than once.

Scope *document, page*

Params See Table 4.14 and Table 4.15.

See also *place_pdi_page()*

Sub *close_pdi_page*(page As Long)

Close the page handle, and free all page-related resources.

page A valid PDF page handle (not a page number!) retrieved with *open_pdi_page()*.

Details This function closes the page associated with the page handle in the PDF document identified by *doc*, and releases all related resources. *page* must not be used after this call.

Scope *document, page*

Params See Table 4.14 and Table 4.15.

Sub *place_pdi_page*(page As Long, x As Single, y As Single, sx As Single, sy As Single)

Place a PDF page with the lower left corner at (x, y), and scale it.

page A valid PDF page handle (not a page number!) retrieved with *open_pdi_page()*. The page handle must not have been closed.

x, y The coordinates in the user coordinate system where the lower left corner of the placed page will be located.

sx, sy The horizontal and vertical scaling factors which will be applied to the page.

Details This function is similar to *place_image()*, but operates on imported PDF pages instead. Another difference is that *place_image()* provides only a single scaling factor, while this function provides separate scaling factors in x and y direction.

Scope *page, pattern, template*

Params See Table 4.14 and Table 4.15.

4.7.2 Parameter Handling

Function *get_pdi_value*(key As String, doc As Long, page As Long, index As Long) As Single

Get some PDI document parameter with numerical type.

key Specifies the name of the parameter to be retrieved, see Table 4.15 and Table 4.14.

doc A valid PDF document handle retrieved with *open_pdi()*.

page A valid PDF page handle (not a page number!) retrieved with `open_pdi_page()`. For keys which are not page-related `page` must be -1.

index Currently unused, must be 0.

Scope *any*

**Function `get_pdi_parameter(`
key As String, doc As Long, page As Long, index As Long) As String

---**

Get some PDI document parameter with string type.

key Specifies the name of the parameter to be retrieved, see Table 4.14 and Table 4.15.

doc A valid PDF document handle retrieved with `open_pdi()`.

page A valid PDF page handle (not a page number!) retrieved with `open_pdi_page()`. For keys which are not page-related `page` must be -1.

index Currently unused, must be 0.

Details This function gets some string parameter related to an imported PDF documented, in some cases further specified by `page` and `index`. Table 4.16 lists relevant parameter combinations.

Scope *any*

Table 4.14. Page-related parameters and values for PDF import

function	key	explanation
<code>get_pdi_value</code>	<code>width</code> <code>height</code>	Get the width or height, respectively, of an imported page in default units. Cropping and rotation will be taken into account.
<code>get_pdi_value</code>	<code>/Rotate</code>	page rotation in degrees (0, 90, 180, or 270)

Table 4.15. Document-related parameters and values for PDF import (the page parameter must be -1)

function	key	explanation
<code>get_parameter</code>	<code>pdi</code>	Returns the string <code>true</code> , if the PDI library is attached (and not restricted to demo mode), and <code>false</code> otherwise.
<code>get_pdi_value</code>	<code>/Root/Pages/Count</code>	total number of pages in the imported document
<code>get_pdi_parameter</code>	<code>filename</code>	name of the PDF file
<code>get_pdi_value</code>	<code>version</code>	PDF version number multiplied by 10, e.g. 13 for PDF 1.3
<code>set_parameter</code>	<code>pdiwarning</code>	This parameter can be used to obtain more detailed information about why a PDF or page couldn't be opened: <code>true</code> Raise a nonfatal exception when the PDI function fails. The information string supplied with the exception may be useful in debugging import-related problems. <code>false</code> Do not raise an exception when the PDI function fails. Instead, the function returns -1 on error. This is default.

Table 4.15. Document-related parameters and values for PDF import (the page parameter must be -1)

function	key	explanation		
set_parameter	pdistrict	This parameter can be used to control PDI's behavior with respect to damaged PDF files:		
		<table border="0"> <tr> <td>true</td> <td>Raise a nonfatal exception for non-conforming PDFs unless the warning parameter is set to false.</td> </tr> <tr> <td>false</td> <td>Accept certain kinds of damaged PDFs. This is the default.</td> </tr> </table>	true	Raise a nonfatal exception for non-conforming PDFs unless the warning parameter is set to false.
true	Raise a nonfatal exception for non-conforming PDFs unless the warning parameter is set to false.			
false	Accept certain kinds of damaged PDFs. This is the default.			

4.8 Hypertext Functions

4.8.1 Document Open Action and Open Mode

Table 4.16 lists relevant parameters and values for this section. These parameters can be set at an arbitrary time before calling `close()`.

Table 4.16. Parameters for document open action and open mode (see Section 4.2.3, »Parameter Handling«)

function	key	explanation
set_parameter	openaction	Set the open action, i.e., the zoom factor for the first page of the document. Possible values are <code>retain</code> , <code>fitpage</code> , <code>fitwidth</code> , <code>fitheight</code> , <code>fitbbox</code> (see Table 4.23). The default is <code>retain</code> . This parameter can be set once at an arbitrary time before <code>close()</code> .
set_parameter	openmode	Set the appearance when the document is opened. The default value is <code>bookmarks</code> if the document contains any bookmarks, and otherwise <code>none</code> :
		<code>none</code> Neither bookmarks nor thumbnails are visible
		<code>bookmarks</code> Open the document with bookmarks visible.
		<code>thumbnails</code> Open document with thumbnails visible
	<code>fullscreen</code> Open the document in fullscreen mode.	

4.8.2 Bookmarks

Table 4.17 lists relevant parameters for this section.

Table 4.17. Parameters for bookmarks (see Section 4.2.3, »Parameter Handling«)

function	key	explanation
set_parameter	bookmark-dest	Set the target zoom for subsequently generated bookmark. Possible values are <code>retain</code> , <code>fitpage</code> , <code>fitwidth</code> , <code>fitheight</code> , <code>fitbbox</code> (see Table 4.23). This parameter can be changed an arbitrary number of times. The default is <code>retain</code> .

Note Adding bookmarks sets the open mode (see Section 4.8.1, »Document Open Action and Open Mode«) to `bookmarks` unless another mode has explicitly been set.

Function `add_bookmark(text As String, parent As Long, open As Long) As Long`

Add a nested bookmark under `parent`, or a new top-level bookmark.

text Contains the text of the bookmark. It may be encoded with `PDFDocEncoding` or `Unicode`. The maximum length of `text` is 255 characters (`PDFDocEncoding`), or 126 Unicode characters. However, a practical limit of 32 characters for `text` is advised.

parent If *parent* contains a valid bookmark handle returned by a previous call to *add_bookmark()*, a new bookmark will be generated which is a subordinate of the given *parent*. In this way, arbitrarily nested bookmarks can be generated. If *parent* = 0 a new top-level bookmark will be generated.

open If 0, child bookmarks will not be visible. If *open* = 1, all children will be folded out. The bookmark target will be viewed at the current bookmark zoom factor which can be set via the *bookmarkdest* parameter (see Table 4.17).

Returns An identifier for the bookmark just generated. This identifier may be used as the *parent* parameter in subsequent calls.

Details This function adds a PDF bookmark with the supplied *text* that points to the current page. The zoom factor can be controlled with the *bookmarkdest* parameter.

Scope *page*

Params *openmode, bookmarkdest*

4.8.3 Document Information Fields

Sub set_info(key As String, value As String)

Fill document information field *key* with *value*.

key The *key* parameter must be encoded with PDFDocEncoding. *key* may be any of the five standard information field names, or an arbitrarily named custom field (see Table 4.18). There is no limit for the number of custom fields. Regarding the use and semantics of custom document information fields, PDFlib users are encouraged to take a look at the Dublin Core Metadata element set.¹

Table 4.18. Values for the document information field *key*

key	explanation
<i>Subject</i>	<i>Subject of the document</i>
<i>Title</i>	<i>Title of the document</i>
<i>Creator</i>	<i>Creator of the document</i>
<i>Author</i>	<i>Author of the document</i>
<i>Keywords</i>	<i>Keywords describing the contents of the document</i>
<i>any name other than CreationDate and Producer</i>	<i>User-defined field. PDFlib supports an arbitrary number of fields. Names consist of printable characters except the following: blank ' ', %, (,), <, >, [,], {, }, /, and #.</i>

value The string to which the *key* parameter will be set. It can be encoded with PDFDocEncoding or Unicode. Acrobat imposes a maximum length of *value* of 255 bytes.

Scope *document, page*

4.8.4 Page Transitions

PDF files may specify a page transition in order to achieve special effects which may be useful for presentations or »slide shows«. In Acrobat, these effects cannot be set docu-

¹ See <http://purl.org/DC>

ment-specific or on a page-by-page basis, but only for the full screen mode. PDFlib, however, allows setting the page transition mode and duration for each page separately. Table 4.19 lists relevant parameters and values for this section.

Table 4.19. Parameters and values for page transitions (see Section 4.2.3, »Parameter Handling«)

function	key	explanation
set_parameter	transition	Set the page transition effect for the current and subsequent pages until the transition is changed again. The transition types below are supported. type may also be empty to reset the transition effect. Default: replace.
	split	Two lines sweeping across the screen reveal the page
	blinds	Multiple lines sweeping across the screen reveal the page
	box	A box reveals the page
	wipe	A single line sweeping across the screen reveals the page
	dissolve	The old page dissolves to reveal the page
	glitter	The dissolve effect moves from one screen edge to another
set_value	duration	Set the page display duration in seconds for the current page. Default: one second.

4.8.5 File Attachments

Sub attach_file(llx As Single, lly As Single, urx As Single, ury As Single, filename As String, description As String, author As String, mimetype As String, icon As String)

Add a file attachment annotation.

llx, lly, urx, ury x and y coordinates of the lower left and upper right corners of the annotation rectangle in default user space coordinates.

filename The name of the file which will be attached to the PDF document.

description A string with some explanation of the attachment. It may be encoded in PDFDocEncoding or Unicode.

author A string with the author's name or function. It may be encoded in PDFDocEncoding or Unicode.

mimetype The MIME type of the file. It will be used by Acrobat for launching the appropriate program when the file attachment annotation is activated.

icon Controls the display of the unopened file attachment in Acrobat (see Table 4.20).

Table 4.20. Icon names for file attachments

icon name	icon appearance	icon name	icon appearance
graph		pushpin	
paperclip		tag	

Details This function adds a file attachment annotation at the specified rectangle. PDF file attachments are only supported in Acrobat 4, and are therefore not supported in PDFlib's Acrobat 3 compatibility mode. Moreover, Acrobat Reader is unable to deal with file at-

tachments and will display a question mark instead. File attachments only work in the full Acrobat software.

Scope page

4.8.6 Note Annotations

Sub `add_note(lly As Single, lly As Single, urx As Single, ury As Single, contents As String, title As String, icon As String, open As Long)`

Add a note annotation.

lly, lly, urx, ury x and y coordinates of the lower left and upper right corners of the note rectangle in default user space coordinates.

contents Text content of the note. It may be encoded with `PDFDocEncoding` or `Unicode`. The maximum length of `contents` is 65535 bytes.

title Heading text of the note. It may be encoded with `PDFDocEncoding` or `Unicode`. The maximum length of `title` is 255 characters (`PDFDocEncoding`), or 126 Unicode characters. However, a practical limit of 32 characters for `title` is advised.

icon Controls the display of the unopened note attachment in Acrobat (see Table 4.21).

open The annotation will be displayed in open state if `open = 1`, and closed if `open = 0`.

Details This function adds a note annotation at the specified rectangle. Different note icons are only available in Acrobat 4, and are not supported in Acrobat 3 compatibility mode (the `icon` parameter must be empty in this case). Acrobat 3 viewers (and apparently Unix versions of Acrobat 4) will display the »note« type icon regardless of the supplied `icon` parameter.

Scope page

Table 4.21. Icon names for note annotations

icon name	icon appearance	icon name	icon appearance
<code>comment</code>		<code>newparagraph</code>	
<code>insert</code>		<code>key</code>	
<code>note</code>		<code>help</code>	
<code>paragraph</code>			

4.8.7 Links

Table 4.22 lists relevant parameters for this section.

Note PDF doesn't support links with shapes other than rectangles.

Table 4.22. Parameters for links (see Section 4.2.3, »Parameter Handling«)

function	key	explanation
<code>set_parameter</code>	<code>base</code>	Set the document's base URL. This is useful when a document with relative Web links to other documents is moved to a different location. Setting the base URL to the »old« location makes sure that relative links will still work.

Sub `add_pdflink(llx As Single, lly As Single, urx As Single, ury As Single, filename As String, page As Long, dest As String)`

Add a file link annotation (to a PDF target).

llx, lly, urx, ury x and y coordinates of the lower left and upper right corners of the link rectangle in default user space coordinates.

filename The name of the target PDF file.

page The physical page number of the target page.

dest The destination zoom. It can attain one of the values specified in Table 4.23.

Scope `page`

Sub `add_loccallink(llx As Single, lly As Single, urx As Single, ury As Single, page As Long, dest As String)`

Add a link annotation to a target within the current PDF file.

llx, lly, urx, ury x and y coordinates of the lower left and upper right corners of the link rectangle in default user space coordinates.

page The physical page number of the target page. This may be a previously generated page, or a page in the same document that will be generated later (after the current page). However, the application must make sure that the target page will actually be generated; PDFlib will issue a warning message otherwise.

dest Specifies the destination zoom, which is one of the values specified in Table 4.23.

Scope `page`

Table 4.23. Values for the `dest` parameter of `add_pdflink()` and `add_loccallink()`. The same values are also used for the `openaction` (see Section 4.8.1, »Document Open Action and Open Mode«) and `bookmarkdest` parameters (see Section 4.8.2, »Bookmarks«).

dest	explanation
<code>retain</code>	Retain the zoom factor which was in effect when the link was activated.
<code>fitpage</code>	Fit the complete page to the window.
<code>fitwidth</code>	Fit the page width to the window.
<code>fitheight</code>	Fit the page height to the window.
<code>fitbbox</code>	Fit the page's bounding box (the smallest rectangle enclosing all objects) to the window.

Sub add_launchlink(lIx As Single, lIy As Single, urx As Single, ury As Single, filename As String)

Add a launch annotation (to a target of arbitrary file type).

lIx, lIy, urx, ury x and y coordinates of the lower left and upper right corners of the link rectangle in default user space coordinates.

filename The name of the file which will be launched upon clicking the link.

Scope page

Sub add_weblink(lIx As Single, lIy As Single, urx As Single, ury As Single, url As String)

Add a weblink annotation to a target URL on the Web.

lIx, lIy, urx, ury x and y coordinates of the lower left and upper right corners of the link rectangle in default user space coordinates.

url A Uniform Resource Identifier encoded in 7-bit ASCII specifying the link target. It can point to an arbitrary (Web or local) resource.

Scope page

Sub set_border_style(style As String, width As Single)

Set the border style for all kinds of annotations.

style Specifies the annotation border style, and must be one of *solid* or *dashed*.

width Specifies the annotation border width in points. If *width = 0* the annotation borders will be invisible.

Details The settings made by this function are used for all annotations until a new style is set. At the beginning of a document the annotation border style is set to a default of a solid line with a width of 1.

Scope document, page

Sub set_border_color(red As Single, green As Single, blue As Single)

Set the border color for all kinds of annotations.

red, green, blue The RGB color values for annotation borders.

Details The settings made by this function are used for all annotations until a new color is set. At the beginning of a document the annotation border color is set to black (0, 0, 0).

Scope document, page

Sub set_border_dash(b As Single, w As Single)

Set the border dash style for all kinds of annotations.

b, w Specify the border dash style (see *setdash()*).

Details At the beginning of a document the annotation border dash style is set to a default of (3, 3). However, this default will only be used when the border style is explicitly set to *dashed*.

Scope document, page

4.8.8 Thumbnails

Sub `add_thumbnail(image As Long)`

Add an existing image as thumbnail for the current page.

image A valid image handle retrieved with one of the `open_image*()` functions, but not a handle to a PDF page or template.

Details This function adds the supplied image as thumbnail image for the current page. A thumbnail image must adhere to the following restrictions:

- ▶ The image must be no larger than 106 x 106 pixels.
- ▶ The image must use the grayscale, RGB, or indexed RGB color space.
- ▶ Multi-strip TIFF images can not be used as thumbnails because thumbnails must be constructed from a single PDF image object, and multi-strip TIFF images result in multiple PDF image objects (see Section 3.4.1, »Supported Image File Formats«).

This function doesn't generate thumbnail images for pages, but only offers a hook for adding existing images as thumbnails. The actual thumbnail images must be generated by the client or some other application. The client must ensure that color, height/width ratio, and actual contents of a thumbnail match the corresponding page contents.

Since Acrobat 5 (including the free Reader) generates thumbnails on the fly, and thumbnails increase the overall file size of the generated PDF, it is recommended not to add thumbnails, but rely on client-side thumbnail generation instead.

Scope page; must only be called once per page. Not all pages must have thumbnails attached to them.

Params `openmode`

4.9 Page Size Formats

For the convenience of PDFlib users, Table 4.24 lists common standard page sizes¹.

Table 4.24. Common standard page size dimensions in points

page format	width	height	page format	width	height
A0	2380	3368	A6	297	421
A1	1684	2380	B5	501	709
A2	1190	1684	letter	612	792
A3	842	1190	legal	612	1008
A4	595	842	ledger	1224	792
A5	421	595	11 x 17	792	1224

¹ More information about ISO, Japanese, and U.S. standard formats can be found at the following URLs: <http://www.twics.com/~eds/papersize.html>, <http://www.cl.cam.ac.uk/~mgk25/iso-paper.html>

5 The PDFlib License

PDFlib is available under two separate licensing terms which are substantially different, and meet the needs of different developer groups. Please take the time to read the short summaries below in order to decide which one applies to your development.

5.1 The »Aladdin Free Public License«

(This section is omitted from this edition of the manual since it does not apply.)

5.2 The Commercial PDFlib License

Different licensing options are available for PDFlib use on one or more servers, and for redistributing PDFlib with your own products. We also offer maintenance and support 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 commercial PDFlib license, or have any questions:

PDFlib GmbH

Tal 40, 80331 München, Germany

<http://www.pdfliib.com>

phone +49 • 89 • 29 16 46 87

fax +49 • 89 • 29 16 46 86

Licensing contact: sales@pdfliib.com

Support for PDFlib licensees: support@pdfliib.com (include your license number)

For other inquiries check our mailing list at <http://www.egroups.com/group/pdfliib>.

6 References

[1] Adobe Systems Incorporated: PDF Reference, Second Edition: Version 1.3. Published by Addison-Wesley 2000, ISBN 0-201-61588-6; also available as PDF from <http://partners.adobe.com/asn/developer/technotes.html>

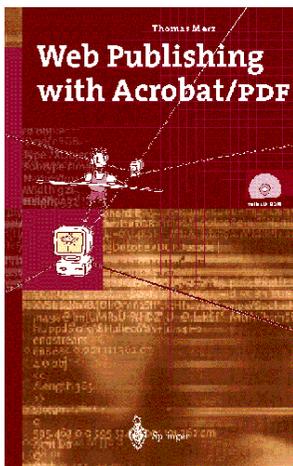
[2] Adobe Systems Incorporated: PostScript Language Reference Manual, third edition. Published by Addison-Wesley 1999, ISBN 0-201-37922-8; also available as PDF from <http://partners.adobe.com/asn/developer/technotes.html>

[3] The following book by the principal author of PDFlib is available in English, German, and Japanese editions. It describes all aspects of integrating PDF in the Web:

English edition: Thomas Merz, *Web Publishing with Acrobat/PDF*. With CD-ROM. Springer-Verlag Heidelberg Berlin New York 1998
ISBN 3-540-63762-1, orders@springer.de
(out of print)

German edition: Thomas Merz, *Mit PDF ins Web*.
Integration, Formulare, Sicherheit, E-Books.
Zweite Auflage. Mit CD-ROM. ISBN 3-935320-00-0, PDFlib Edition 2001
PDFlib GmbH, 80331 München, Tal 40, fax +49 • 89 • 29 16 46 86
<http://www.pdflib.com>, books@pdflib.com

Japanese edition: Tokyo Denki Daigaku 1999, ISBN 4-501-53020-0
<http://plaza4.mbn.or.jp/~unit>



A Shared Libraries and DLLs

(This section is not included in this edition of the PDFlib manual.)

B PDFlib Quick Reference

General Functions

Function prototype	page
<i>Function open_file(filename As String) As Long</i>	71
<i>Function get_buffer()</i>	71
<i>Sub close()</i>	72
<i>Sub begin_page(width As Single, height As Single)</i>	72
<i>Sub end_page()</i>	72
<i>Function get_value(key As String, modifier As Single) As Single</i>	73
<i>Sub set_value(key As String, value As Single)</i>	73
<i>Function get_parameter(key As String, modifier As Single) As String</i>	73
<i>Sub set_parameter(key As String, value As String)</i>	74

Text Functions

Function prototype	page
<i>Function findfont(fontname As String, encoding As String, embed As Long) As Long</i>	74
<i>Sub setfont(font As Long, fontsize As Single)</i>	75
<i>Sub show(text As String)</i>	76
<i>Sub show_xy(text As String, x As Single, y As Single)</i>	77
<i>Sub continue_text(text As String)</i>	77
<i>Function show_boxed(text As String, x As Single, y As Single, width As Single, height As Single, mode As String, feature As String) As Long</i>	77
<i>Function stringwidth(text As String, font As Long, size As Single) As Single</i>	78
<i>Sub set_text_pos(x As Single, y As Single)</i>	78

Graphics Functions

Function prototype	page
<i>Sub setdash(b As Single, w As Single)</i>	79
<i>Sub setpolydash(darray)</i>	79
<i>Sub setflat(flatness As Single)</i>	79
<i>Sub setlinejoin(linejoin As Long)</i>	80
<i>Sub setlinecap(linecap As Long)</i>	80
<i>Sub setmiterlimit(miter As Single)</i>	81
<i>Sub setlinewidth(width As Single)</i>	81
<i>Sub save()</i>	82
<i>Sub restore()</i>	82
<i>Sub translate(tx As Single, ty As Single)</i>	82
<i>Sub scale(sx As Single, sy As Single)</i>	82
<i>Sub rotate(phi As Single)</i>	83
<i>Sub skew(alpha As Single, beta As Single)</i>	83
<i>Sub concat(a As Single, b As Single, c As Single, d As Single, e As Single, f As Single)</i>	83

Function prototype	page
<i>Sub moveto(x As Single, y As Single)</i>	84
<i>Sub lineto(x As Single, y As Single)</i>	84
<i>Sub curveto(x1 As Single, y1 As Single, x2 As Single, y2 As Single, x3 As Single, y3 As Single)</i>	85
<i>Sub circle(x As Single, y As Single, r As Single)</i>	85
<i>Sub arc(x As Single, y As Single, r As Single, alpha As Single, beta As Single)</i>	85
<i>Sub arcn(x As Single, y As Single, r As Single, alpha As Single, beta As Single)</i>	86
<i>Sub rect(x As Single, y As Single, width As Single, height As Single)</i>	86
<i>Sub closepath()</i>	86
<i>Sub stroke()</i>	87
<i>Sub closepath_stroke()</i>	87
<i>Sub fill()</i>	87
<i>Sub fill_stroke()</i>	87
<i>Sub closepath_fill_stroke()</i>	87
<i>Sub clip()</i>	88
<i>Sub endpath()</i>	88

Color Functions

Function prototype	page
<i>Sub setcolor(type As String, colorspace as String, c1 As Single, c2 As Single, c3 As Single, c4 As Single)</i>	88
<i>Function makespotcolor(spotname As String) As Long</i>	89
<i>Function begin_pattern(width As Single, height As Single, xstep As Single, ystep As Single, painttype As Long) As Long</i>	89
<i>Sub end_pattern(pattern As Long)</i>	90

Image Functions

Function prototype	page
<i>Function open_image_file(type As String, filename As String, stringparam As String, intparam As Long) As Long</i>	91
<i>Function open_CCITT(filename As String, width As Long, height As Long, BitReverse As Long, K As Long, BlackIs1 As Long) As Long</i>	91
<i>Function open_image(type As String, source As String, data, length As Long, width As Long, height As Long, components As Long, bpc As Long, params As String) As Long</i>	92
<i>Sub close_image(image As Long)</i>	93
<i>Sub place_image(image As Long, x As Single, y As Single, scale As Single)</i>	94
<i>Function begin_template(width As Long, height As Long) As Long</i>	94
<i>Sub end_template()</i>	94

PDF Import (PDI) Functions

Function prototype	page
Function <i>open_pdi</i> (filename As String, stringparam As String, intparam As Long) As Long	95
Sub <i>close_pdi</i> (doc As Long)	95
Function <i>open_pdi_page</i> (doc As Long, pagenumber As Long, pagelabel As String) As Long	95
Sub <i>close_pdi_page</i> (page As Long)	96
Sub <i>place_pdi_page</i> (page As Long, x As Single, y As Single, sx As Single, sy As Single)	96
Function <i>get_pdi_value</i> (key As String, doc As Long, page As Long, index As Long) As Single	96
Function <i>get_pdi_parameter</i> (key As String, doc As Long, page As Long, index As Long) As String	97

Hypertext Functions

Function prototype	page
Function <i>add_bookmark</i> (text As String, parent As Long, open As Long) As Long	98
Sub <i>set_info</i> (key As String, value As String)	99
Sub <i>attach_file</i> (llx As Single, lly As Single, urx As Single, ury As Single, filename As String, description As String, author As String, mimetype As String, icon As String)	100
Sub <i>add_note</i> (llx As Single, lly As Single, urx As Single, ury As Single, contents As String, title As String, icon As String, open As Long)	101
Sub <i>add_pdflink</i> (llx As Single, lly As Single, urx As Single, ury As Single, filename As String, page As Long, dest As String)	102
Sub <i>add_locallink</i> (llx As Single, lly As Single, urx As Single, ury As Single, page As Long, dest As String)	102
Sub <i>add_launchlink</i> (llx As Single, lly As Single, urx As Single, ury As Single, filename As String)	103
Sub <i>add_weblink</i> (llx As Single, lly As Single, urx As Single, ury As Single, url As String)	103
Sub <i>set_border_style</i> (style As String, width As Single)	103
Sub <i>set_border_color</i> (red As Single, green As Single, blue As Single)	103
Sub <i>set_border_dash</i> (b As Single, w As Single)	103
Sub <i>add_thumbnail</i> (image As Long)	104

Parameters and Values

category	function	keys
setup	<i>set_parameter</i>	prefix, resourcefile, compatibility, serial, warning
	<i>set_value</i>	compress, floatdigits
	<i>get_value</i>	major, minor, revision
	<i>get_parameter</i>	version
document	<i>set_value</i>	pagewidth, pageheight CropBox, BleedBox, ArtBox, TrimBox: these must be followed by a slash '/' character and one of llx, lly, urx, ury, for example: CropBox/llx
font	<i>set_parameter</i>	FontAFM, FontPFM, FontOutline, Encoding, fontwarning
text	<i>set_value</i>	leading, textrise, horizscaling, textrendering, charspacing, wordspacing
	<i>get_value</i>	leading, textrise, horizscaling, textrendering, charspacing, wordspacing, textx, texty, font, fontsize, capheight, ascender, descender
	<i>set_parameter</i>	underline, overline, strikeout, nativeunicode
	<i>get_parameter</i>	underline, overline, strikeout, fontname, fontencoding

category	function	keys
graphics	<i>set_parameter</i>	<i>fillrule</i>
	<i>get_value</i>	<i>currentx, currenty</i>
image	<i>get_value</i>	<i>imagewidth, imageheight, resx, resy</i>
	<i>set_parameter</i>	<i>imagewarning</i>
PDI	<i>get_parameter</i>	<i>pdi</i>
	<i>set_parameter</i>	<i>pdiwarning, pdistrict</i>
	<i>get_pdi_value</i>	<i>/Root/Pages/Count, /Rotate, version, width, height</i>
hypertext	<i>get_pdi_parameter</i>	<i>filename</i>
	<i>set_parameter</i>	<i>openaction, openmode, bookmarkdest, transition, base</i>
	<i>set_value</i>	<i>duration</i>

C Revision History

Version information on PDFlib can be found in the source distribution.

Revision history of this manual

Date	Changes
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 ActiveX 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

O-9

16-bit encoding 52
8-bit encodings 37, 52

A

Acrobat 3 compatibility 10
Acrobat 4 compatibility 10
Acrobat 5 compatibility 11
Active Server Pages 18
 special considerations 18
ActiveX binding
 error handling 17
 general 14
 redistribution 16
 register the DLL 16
 Unicode support 18
 version control 18
add_bookmark() 98
add_launchlink() 103
add_locallink() 102
add_note() 101
add_pdflink() 102
add_thumbnail() 104
add_weblink() 103
Adobe Font Metrics (AFM) 42
Adobe Glyph List (AGL) 40, 52
AFM (Adobe Font Metrics) 42
AGL (Adobe Glyph List) 40, 52
Aladdin free public license 105
All spot color name 89
Allaire ColdFusion: *see* ColdFusion
alpha channel 62
annotations 52, 101
API (Application Programming Interface)
 reference 69
arc() 85
arcn() 86
ArtBox 33, 71, 110
AS/400 37
ascender 54
ascender parameter 74
Asian FontPack 48
attach_file() 100
attachments 52, 100
Author field 99
automation 15
availability of PDFlib 12

B

base parameter 102
baseline compression 58
begin_page() 72
begin_pattern() 89
begin_template() 94
Bézier curve 85
bindings 12
BitReverse 92
Black1st 92
BleedBox 33, 71, 110
blind mode 56, 78
bold CJK text 51
bookmarkdest parameter 98
bookmarks 52, 98
 hide 98
Borland Delphi: *see* Delphi
builtin encoding 37, 38

C

C binding 28
C++ binding 28
capheight 54
capheight parameter 74
categories of resources 46
CCITT 59, 91
character ID (CID) 48
character metrics 54
character names 41
character sets 37, 52
characters per inch 55
charspacing parameter 76
Chinese 48, 50
CID fonts 48
circle() 85
 Problems with VB 23
CJK (Chinese, Japanese, Korean) 48
clip 34
clip() 88
close() 72
close_image() 93
close_pdi() 95
close_pdi_page() 96
closepath() 86
closepath_fill_stroke() 87
closepath_stroke() 87
CMaps 48, 50
code page
 8-bit 38

- IBM 1047 38
- Microsoft Windows 1250-1258 40
- Microsoft Windows 1252 37
- Unicode-based 52
- ColdFusion 25
 - error handling 26
 - special considerations 25
 - Unicode support 26
- color 34
- color functions 88
- COM (Component Object Model): see ActiveX
 - binding
- commercial license 105
- compatibility
 - Acrobat 4 10
 - Acrobat Reader 10
- compatibility parameter 70
- compress parameter 71
- concat() 83
- continue_text() 77
- coordinate range 33
- coordinate system 31, 81
 - metric 31
 - top-down 32
- core fonts 36
- CPI (characters per inch) 55
- Creator field 99
- CropBox 33, 71, 110
- current point 34
- currentx and currenty parameter 84
- curveto() 85
- custom encoding 38

D

- default coordinate system 31
- default zoom 98
- Delphi 26
 - error handling 27
 - special considerations 26
 - Unicode support 27
- demo stamp 29
- descender 54
- descender parameter 74
- descriptor 41
- DLL (dynamic link library) 17, 107
- document and page functions 71
- document information fields 52, 99
- document open action 98
- downsampling 60
- dpi calculations 60
- Dublin Core 99
- duration parameter 100

E

- EBCDIC 14, 38
- ebcdic encoding 37, 38
- embedded systems 12

- embedding fonts 41
- encoding 37, 52
 - CJK 51
 - custom 38
 - for hypertext 41
 - Unicode 53
- Encoding parameter 74
- end_page() 72
- end_pattern() 90
- end_template() 94
- endpath() 88
- error handling 12, 30
 - in ActiveX 17
 - in ColdFusion 26
 - in Delphi 27
 - in JScript 21
 - in VBScript 21
 - in Visual Basic 23
- Errors class 18
- Euro character 37, 42
- evaluation stamp 29
- exception: see error handling
- external image references 62

F

- features of PDFlib 9
- file attachments 52, 100
- filename parameter 97
- fill 34
- fill() 87
- fill_stroke() 87
- fillrule parameter 86
- findfont() 74
- font metrics 54
- font parameter 74
- FontAFM parameter 74
- fontencoding parameter 74
- fontname parameter 74
- FontOutline parameter 74
- FontPFM parameter 74
- fonts
 - AFM files 42
 - Asian fontpack 48
 - CID fonts 48
 - CJK fonts 48
 - descriptor 41
 - embedding 41
 - general 36
 - glyph names 41
 - legal aspects of embedding 45
 - monospaced 55
 - OpenType 44
 - PDF core set 36
 - PFA files 43
 - PFB files 43
 - PFM files 42
 - PostScript 42, 43
 - resource configuration 45

- TrueType 44
- type 1 43
- type 3 43
- Unicode support 52
- fontsize parameter 74
- FontSpecific encoding 38
- fontwarning parameter 74
- form XObjects 35

G

- general graphics state 79
- get_buffer() 30, 71
- get_parameter() 73
- get_pdi_parameter() 96, 97
- GIF 58, 91
- global.asa 19
- graphics functions 79
- graphics state 79, 81
- grid.pdf 31
- GUID 18

H

- hello world example
 - for ASP/JScript 19
 - for ASP/VBScript 20
 - general 12
 - in ColdFusion 25
 - in Delphi 26
 - in VBScript 24
 - in Visual Basic 23
- horizontal writing mode 48, 51
- horzscaling parameter 76
- host encoding 37, 38
- host fonts 44
- hypertext
 - encoding 41
 - functions 98

I

- icons
 - for file attachments 100
 - for notes 101
- ignoremask 64, 91
- image data, re-using 61
- image file formats 58
- image functions 90
- image mask 62
- image references 62
- image scaling 59, 60
- imagewarning parameter 58, 90
- imagewidth and imageheight parameters 90
- import functions for PDF 95
- inch 31
- in-core PDF generation 29
- initgraphics() 81
- installation, silent for ActiveX 16

- InstallShield 16
- Internet Service Provider 17
- invert 91
- invisible text 56
- ISO 10646 52
- ISO 8859-1 37, 41
- ISO 8859-2 to -15 40
- ISO 8859-2 39

J

- Japanese 48, 50
- Java binding 28
- JPEG 10, 58, 91, 93
- JScript 19
 - error handling 21
 - Unicode support 22

K

- K parameter for CCITT images 92
- Keywords field 99
- Korean 48, 50

L

- landscape mode 72
- language bindings: see bindings
- Latin 1 encoding 37, 41
- Latin 2 encoding 39
- leading 54
- leading parameter 76
- licensing conditions 105
- licensing PDFlib and PDI 29
- line spacing 54
- lineto() 84
- links 101
- LZW compression 58, 68

M

- macroman encoding 37
- major parameter 71
- makepsres utility 45
- makespotcolor() 89
- mask 63, 91
- masked 63, 91
- masking images 62
- MediaBox 33
- memory images 62
- memory, generating PDF documents in 29
- metadata 99
- metric coordinates 31
- metrics 54
- Microsoft Transaction Server 15
- millimeters 31
- minor parameter 71
- mirroring 83
- monospaced fonts 55
- moveto() 84

MTS 15
multi-page image files 64

N

nagger 29
nativeunicode parameter 54, 76
None spot color name 89
non-proportional image scaling 61
note annotations 52, 101

O

open_CCITT() 91
open_file() 71
open_image() 92
open_image_file() 91
open_pdi() 95
open_pdi_page() 95
openaction parameter 98
openmode parameter 98
OpenType fonts 44
outline text 56
output accuracy 33
overline parameter 55, 76

P

page 64
page descriptions 31
page parameter 91
page size formats 104
 limitations in Acrobat 33
page transitions 99
pagewidth and pageheight parameters 71
parameter
 ascender 74
 base 102
 bookmarkdest 98
 capheight 74
 charspacing 76
 compatibility 70
 compress 71
 currentx and currenty 84
 descender 74
 duration 100
 Encoding 74
 filename 97
 fillrule 86
 font 74
 FontAFM 74
 fontencoding 74
 fontname 74
 FontOutline 74
 FontPFM 74
 fontsize 74
 fontwarning 74
 horizscaling 76
 imagewarning 58, 90

 imagewidth and imageheight 90
 leading 76
 major 71
 minor 71
 nativeunicode 54, 76
 openaction 98
 openmode 98
 overline 55, 76
 pageheight and pagewidth 71
 pdi 97
 pddistrict 98
 pdiwarning 97
 prefix 70
 resourcefile 70
 resx and resy 90
 revision 71
 serial 70
 strikeout 55, 76
 textrendering 51, 56, 76
 textrise 76
 textx and texty 57, 76
 transition 100
 underline 55, 76
 version 71, 97
 warning 31, 70
 width and height 97
 wordspacing 76
parameter handling functions 73
path 34
 painting and clipping 86
pattern color space 35
PDF import functions 95
PDF import library (PDI) 65
PDFDocEncoding 41
PDFlib
 features 9
 program structure 29
 thread-safety 9, 14
PDI 65
pdi parameter 97
pddistrict parameter 98
pdiwarning parameter 97
Perl binding 28
PFA (Printer Font ASCII) 43
PFB (Printer Font Binary) 43
PFM (Printer Font Metrics) 42
Photoshop 93
PHP binding 28
place_image() 94
place_pdi_page() 96
platforms 12
PNG 58, 63, 91
Portable Document Format Reference Manual 106
PostScript fonts 42, 43
PostScript Language Reference Manual 106
prefix parameter 70
print_glyphs.ps 41

Printer Font ASCII (PFA) 43
Printer Font Binary (PFB) 43
Printer Font Metrics (PFM) 42
ProgID 18
program structure 29
Python binding 28

R

raster images
 functions 90
 general 58
raw image data 59, 92
rect() 86
redistributing the ActiveX component 16
references 106
reflection 83
register the PDFlib ActiveX DLL 16
registry 18
regsvr32 17
resource category 46
resourcefile parameter 70
restore() 82
resx and resy parameter 90
revision parameter 71
RGB color 34
rotate() 83
rotating objects 32
RPG binding 28

S

S/390 37
save() 82
scale() 82
 problems with VB 23
scaling images 59, 60
scope definitions 69
separation color space 34
serial parameter 70
serial string 29
set_border_color() 103
set_border_dash() 103
set_border_style() 103
set_info() 99
set_text_pos() 78
set_value() 73
setcolor() 88
setdash() 79
setflat() 79
setfont() 75
setgray() 90
setgray_fill() 90
setgray_stroke() 90
setlinecap() 80
setlinejoin() 80
setlinewidth() 81
setmatrix() 84

setmiterlimit() 81
setpolydash() 79
setrgbcolor() 90
setrgbcolor_fill() 90
setrgbcolor_stroke() 90
setup functions 70
shared libraries 107
show() 76
show_boxed() 56, 77
show_xy() 77
silent install 16
skew() 83
special graphics state 81
spot color (separation color space) 34
standard page sizes 104
strikeout parameter 55, 76
stringwidth() 56, 78
stroke 34
stroke() 87
structure of PDFlib programs 29
Subject field 99
subpath 34
subscript 55, 76
superscript 55, 76
Symbol font 38

T

T1lib 43
Tcl binding 28
templates 35
text box formatting 54
text functions 74
text handling 36
text metrics 54
text rendering modes 56
text variations 54
textrendering parameter 51, 56, 76
textrise parameter 76
textx and texty parameter 51, 57, 76
threading model 15
thread-safety 9, 14
thumbnails 98, 104
TIFF 59, 91
 multi-page 64
tiling 35
Title field 99
top-down coordinates 32
transition parameter 100
translate() 82
transparency 62
 problems with 64
TrimBox 33, 71, 110
TrueType fonts 44
TTF (TrueType font) 44
type 1 fonts 43
type 3 fonts 43

U

- UNC 18
- underline parameter 55, 76
- Unicode 14, 52
 - in ActiveX 18
 - in ColdFusion 26
 - in Delphi 27
 - in JScript 22
 - in VBScript 22
 - in Visual Basic 24
 - problems with language bindings 53
- units 31
- UPR (Unix PostScript Resource) 45
 - file format 46
 - file searching 47
- URL 62, 103
- user space 31

V

- value: *see* parameter
- VBA 15
- vbObjectError 18
- VBScript 20, 24
 - error handling 21
 - Unicode support 22
- version control
 - general 13
 - in ActiveX 18

- version parameter 71, 97
- vertical writing mode 48, 51
- virtual directory 18
- Visual Basic 22
 - error handling 23
 - special considerations 22
 - Unicode support 24
- Visual Basic for Applications 15
- Visual Basic Scripting Edition: *see* VBScript

W

- warning 33, 102
- warning parameter 31, 70
- weblink 103
- width and height parameters 97
- winansi encoding 37
- Windows registry 18
- Windows Script Host (WSH) 24
- wordspacing parameter 76
- writing modes 48, 51

X

- XObjects 35

Z

- ZapfDingbats font 38
- zoom factor 98